

LOAN DOCUMENT

PHOTOGRAPH THIS SHEET


①

INVENTORY

AD-A268 426



DTIC ACCESSION NUMBER



LEVEL

AFOSR-TR-93-0493

DOCUMENT IDENTIFICATION

25 Jan 91

DISTRIBUTION STATEMENT A

Approved for public release:
Distribution Unlimited

DISTRIBUTION STATEMENT

ACCESSION FOR	
NTIS	GRAB <input checked="" type="checkbox"/>
DTIC	TRAC <input type="checkbox"/>
UNANNOUNCED	<input type="checkbox"/>
JUSTIFICATION	
BY	
DISTRIBUTION/	
AVAILABILITY CODES	
DISTRIBUTION	AVAILABILITY AND/OR SPECIAL
A-1	

DISTRIBUTION STAMP

DTIC QUALITY INSPECTED 3

DTIC
ELECTE
AUG 12 1993
S C D

DATE ACCESSIONED

DATE RETURNED

93-18605



59106

REGISTERED OR CERTIFIED NUMBER

93 8 10 227

DATE RECEIVED IN DTIC

PHOTOGRAPH THIS SHEET AND RETURN TO DTIC-FDAC

H
A
N
D
L
E
W
I
T
H
C
A
R
E

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED			1b. RESTRICTIVE MARKINGS	
2a. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION/AVAILABILITY OF REPORT APPROVED FOR PUBLIC RELEASE: DISTRIBUTION UNLIMITED	
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE				
4. PERFORMING ORGANIZATION REPORT NUMBER(S) AFOSR/ONR FINAL TECHNICAL REPORT			5. MONITORING ORGANIZATION REPORT NUMBER(S) AFOSR-TR 88 0012	
6a. NAME OF PERFORMING ORGANIZATION HARVARD UNIVERSITY	6b. OFFICE SYMBOL (if applicable)	7a. NAME OF MONITORING ORGANIZATION Air Force Office of Scientific Research		
6c. ADDRESS (City, State, and ZIP Code) 1236 William James Hall 33 Kirkland Street Cambridge, MA 02138-2044		7b. ADDRESS (City, State, and ZIP Code) Building 410 Bolling Air Force Base, D.C. 20332-6448		
8a. NAME OF FUNDING/SPONSORING ORGANIZATION AFOSR	8b. OFFICE SYMBOL (if applicable) NL	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER AFOSR-12 88-0012		
8c. ADDRESS (City, State, and ZIP Code) Building 410 Bolling Air Force Base, D.C. 20332-6448		10. SOURCE OF FUNDING NUMBERS		
		PROGRAM ELEMENT NO. 61102F	PROJECT NO. 2313	TASK NO. A4
11. TITLE (Include Security Classification) The Neuropsychology of Imagery Processing				
12. PERSONAL AUTHOR(S) Stephen M. Kosslyn				
13a. TYPE OF REPORT Final Tech. Report	13b. TIME COVERED FROM 12/87 TO 11/90	14. DATE OF REPORT (Year, Month, Day) 1991, January, 25	15. PAGE COUNT 58	
16. SUPPLEMENTARY NOTATION				
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number) Neuropsychology, vision, computational models, spatial orientation	
FIELD	GROUP	SUB-GROUP		
05	09			
19. ABSTRACT (Continue on reverse if necessary and identify by block number) High-level visual processes make use of stored information, and are invoked during object identification, navigation, tracking, and visual mental imagery. The present work has resulted in a theory of the component "processing subsystems" used in high-level vision. This theory was developed by considering neuroanatomical, neurophysiological, and computational constraints. The theory has led to two kinds of empirical work: First, specific hypotheses about individual processing subsystems have been tested. For example, the analysis of the representation of spatial relations led to the prediction that two subsystems are used to encode this information, and a set of experiments was conducted that provided support for this distinction. This work has involved a combination of divided-visual-field experiments with normal subjects and detailed examinations of patients with focal brain damage. Second, the subsystems have been implemented in a large computer simulation model, which has been used to generate predictions about specific neurological syndromes. The model can be damaged in a variety of ways, and its performance on a set of tasks then observed. Predictions from the model have been tested, the the results generally support its underlying assumptions and specific claims. In addition, individual subsystems have been implemented as "neural network" simulation models, which have been related directly to properties of the neural substrate assumed to underlie processing. The experiments conducted to date are summarized in the context of the theory in this report, and the utility of the theory for understanding the effects of brain damage is illustrated by reviewing a single case study in detail.				
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input checked="" type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED	
22a. NAME OF RESPONSIBLE INDIVIDUAL Alfred R. Freely			22b. TELEPHONE (Include Area Code) (202) 767-5021	22c. OFFICE SYMBOL NL

Final Technical Report
Neuropsychological Component
of Imagery Processing

AEOSR-TA 93 0493

AFOSR Grant 88-0012
 S. M. Kosslyn, PI

Visual mental imagery is one of the few cognitive abilities that can be easily related to brain function. It has been shown convincingly that visual mental imagery shares mechanisms with visual perception (e.g., for a review see Farah, 1988), and we know an enormous amount about the neural substrate of vision. In addition, imagery clearly relies on memory, and we also know a lot about the neural mechanisms underlying memory (e.g., Squire, 1987). One reason we know so much about vision and memory is that nonhuman primates have similar systems, and so animal models can be studied to understand these abilities. Animal models are not available for many other cognitive abilities, such as language. In this chapter I outline some ways in which findings about the neural substrates of vision and memory can inspire theories of human visual mental imagery.

Two kinds of work have progressed over the course of this grant, theoretical and empirical. Rather than summarize each individual paper, which have been cited in the Annual Reports, I will synthesize what we have learned. In the first part of this report, I will summarize the theoretical inferences we have drawn, and will briefly cite some of the relevant findings that have led to these inferences. In the second part, I will illustrate how we have used these inferences to study patients who have suffered brain damage. For illustrative purposes, I will describe our detailed study of a single patient. We have studied 9 patients in detail to date, with each leading us to different inferences about underlying processing.

I. A Theory of Visual Cognition

Over the course of this grant, my colleagues and I (e.g., Kosslyn, 1987; Kosslyn, Flynn, Amsterdam, & Wang, 1990), have developed a theory of visual cognition. We have used an approach that relies not only on results from neuroanatomy and neurophysiology, but also on computational analyses of how a machine with the structure of the brain could function in specific ways. Before beginning, then, I must briefly outline some key properties of imagery that must be explained. Following this, I will consider the implications of facts about the primate visual system and memory system for how the brain might produce these behaviors.

Key Phenomena to be Explained

Visual mental imagery is a complex phenomenon that has many distinct facets. We have focused on behaviors that reflect the nature, formation, and use of image representations.

Geometric representation

Visual imagery is used to help one recall information about previously perceived objects and events, to reason about visual and spatial properties of objects, and to learn new information (see Kosslyn, Segar, Hillger, & Pani, in press). In all of these circumstances, the local geometry of surfaces of objects must be made explicit. Kosslyn (1980) argued that an array representation is an efficient way of serving this end. If images are patterns of points in a short-term memory structure that functions as an array, the spatial relations among portions of an object are depicted.

Generation

One of the most obvious facts about visual mental imagery is that we do not have images all of the time. Images come and go, depending on the situation. Patterns in the array are best viewed as short-term memory representations. Thus, there must be means of both storing visual representations in long-term memory, and activating the representations to form images in the array.

Part of our ability to activate images involves combining images of different objects into novel combinations. For example, one can imagine Margaret Thatcher riding a zebra, and determine whether she could see over the top of the zebra's head. Indeed, much of the power of imagery comes not only from the ability to image new combinations of objects, but also from the ability to generate new patterns; one can "mentally draw" in imagery, producing images of patterns never actually seen.

Inspection

Patterns in an array would be useless if they could not be interpreted. For example, if one is asked to image an upper case letter "a" and then to mentally rotate it 180°, most people can report the shape of the enclosed area (a triangle balanced on its apex). We must have some way of interpreting the patterns in images. Furthermore, we can "zoom in" on isolated parts of imaged patterns or scan across them (see Kosslyn, 1980, for reviews of experiments demonstrating these abilities).

Recoding

Not only can we interpret patterns in images, but we also can encode them into memory (cf. Paivio, 1971). After imagining objects in new combinations, or imaging new patterns altogether, we can remember them.

Maintenance

Many of our imagery abilities are limited by the fact that images require effort to maintain. The more perceptual units that are included in an image, the more difficult it is to maintain (see Kirby & Kosslyn, in press; Kosslyn, 1980).

Transformation

Finally, the ability to transform imaged patterns lies at the heart of the use of imagery in reasoning. For example, we can rotate patterns in images, including in the third dimension so that we "see" new portions as they come into view. We also can imagine objects growing or shrinking (Shepard & Cooper, 1982), and probably can perform many other types of transformations as well.

Any theory of imagery must provide accounts for these basic properties. The continued development of a theory of imagery in our laboratory is driven by this requirement. We have found numerous insights into these phenomena by considering facts about the brain, as is discussed in the following section.

A Cognitive Neuroscience of Imagery Processing

Kosslyn, Flynn, Amsterdam, and Wang (1990) described a theory of visual object identification. This theory posits a set of processing subsystems that work together to identify shapes and specify their locations. A processing subsystem corresponds to a neural network or set of related neural networks (i.e., which work together to perform part of an information processing task) and is defined by the type of input it accepts, the operation it performs on the input, and the type of output it produces (which in turn serves as input to other subsystems).

Kosslyn (1987) used an early version of the Kosslyn et al. theory to understand the relationship between visual mental imagery and visual perception, which has since been carried further by Kosslyn (in press). Our key assumption is that visual mental imagery shares processing subsystems with visual perception, which seems reasonable given the confluence of findings from numerous experiments using various methodologies (see Farah, 1988).

In this section I briefly describe each subsystem posited by the Kosslyn et al. theory, as well as how the subsystems are interconnected. In each case, I will describe the role of a subsystem in vision before turning to imagery, and will note the ways in which the previous theory has been modified. The architecture of the system underlying visual object recognition and identification is illustrated in Figure 1.

Insert Figure 1 About Here

Input to the system

The input to high-level vision is a representational structure that stores the output from low-level visual processes in perception (i.e., those driven purely by stimulus input, which detect edges, color, and so on); selected contents of this structure are then passed on for further processing.

Visual buffer. High-level visual processes take as input the patterns of activation in a series of topographically mapped areas of cortex. There are at least 15 such maps in the primate brain (for recent reviews, see Felleman & Van Essen, in press, and Van Essen, Felleman, DeYoe, Olavarria, & Knierim, 1990). I focus on the topographically mapped areas following V1 (and perhaps V2) in the processing stream (V1 apparently is dedicated to low-level visual processing), and conceive of these structures as forming a single functional structure that I call the *visual buffer*. The areas subsumed by this structure are localized in circumstriate cortex in the occipital lobe.

The visual buffer corresponds to the array in the theory of Kosslyn (1980). Kosslyn (1987) noted that the topographically mapped areas of cortex receive connections not only from the lower visual areas, but also from the higher ones. Thus, it is possible that a visual mental image is a pattern of activation in the visual buffer that is induced by stored information, as opposed to input from the eyes (which induces a pattern of activation during perception).

Kosslyn (1980) treated the visual buffer as a static structure, exactly analogous to an array in a computer. This seems overly simplistic. My present view is that the visual buffer itself performs much computation. I suspect that we do not store very complete information in long-term memory, and that when an image is generated the buffer itself must fill in many gaps in patterns. This filling-in process may rely on bottom-up processes that complete fragments that are colinear, fill in regions of the same color, texture, and so forth. This sort of processing would allow stored fragments to engender a more complete pattern.

If some of the topographically mapped areas used in perception are also used in imagery, then at least some of the limits on our ability to maintain visual mental images make sense: In perception, one does not want smearing as one moves one's eyes from place to place. Thus the visual buffer does not retain patterns of activation long. This property is inherited in imagery, which uses the same structure—and so images fade quickly and require effort to maintain.

Furthermore, another property of the topographically mapped cortical areas allows us to understand why individual parts are hard to "see" when an object is imaged at a small size. "Spatial summation" is a neural averaging over variations within a given region, and is common within these visual areas. This property would also affect images, introducing a "grain" to the array; if objects are too small (i.e., cover too small a region of the visual buffer), details will not be represented.

Attention window. The visual buffer typically contains more information than can be processed during perception (there are more cells in these areas than there are projections to other visual areas; cf. Van Essen, 1985). Hence, some information must be given a high priority for further processing whereas other information must be placed in the background. The *attention window* selects a region within the visual buffer for detailed further processing. The size of the window in the visual buffer can be altered (cf. Larsen & Bundesen, 1978; Treisman & Gelade, 1980). Indeed, Larsen and Bundesen (1978) and Cave and Kosslyn (1989) showed that the time necessary to adjust the size of the attention window increases linearly with the amount of adjustment necessary.

In addition, the location of the attention window in the visual buffer can be shifted, independently of any overt attention shift. Kosslyn (1973) showed that people can scan visual mental images, even when their eyes are closed, and the farther they scan across the imaged object, the more time is required.

However, we do not "bump into the edge" of the visual buffer when we scan; rather, we can scan to portions of objects that initially were "off screen" (see Kosslyn, 1980, for evidence). This can be accomplished if new portions of an image are introduced on one side of the visual buffer and the pattern is slid towards the opposite side (rather like an image on a TV screen as the camera scans over a scene). Similarly, when we "zoom in" on an imaged object, further details of the object become apparent. Thus, there may be a means of fixing a portion of a pattern in the attention window, and adding more details to the pattern as the window is expanded.

Subsystems of the ventral system

A major anatomical pathway runs from the occipital lobes down to the inferior temporal lobes, which has been shown to be involved in the representation of object properties such as shape and color (e.g., Maunsell & Newsome, 1987; Mishkin, Ungerleider, & Macko, 1983; and Ungerleider & Mishkin, 1982). This "ventral system" receives the information that is selected by the attention window. Kosslyn et al. (1990) decompose the ventral system into three subsystems.

Preprocessing. A vision system must be able to produce the same perceptual representation for an object when it is viewed in different locations in the visual field and from different points of view. Whenever a range of different inputs must be mapped to the same output, one seeks a set of common properties (or overlapping properties, exploiting Wittgensteinian "family resemblances"). Lowe (1987a, 1987b) calls these "nonaccidental properties" (see also Biederman, 1987). For example, properties such as parallel lines (usually indicating edges), line intersections, and symmetries are likely to remain invariant under translation, rotation, and scale changes. Some subsystem presumably

computes these useful invariants for subsequent matching against stored information. Not all of the properties are likely to be preserved for all objects, but one cannot know that until the object has been identified; thus, the subsystem must operate in large part purely on the basis of the stimulus input. Kosslyn et al. (1990) hypothesize that such a preprocessing subsystem is implemented in the occipital-temporal area, which receives information from the lower visual areas in the occipital lobes, and sends information to higher visual areas in the temporal lobes.

Lowe's conception of nonaccidental properties is very powerful in certain domains, such as recognizing many manufactured objects. However, many natural objects are not easily described using such properties (e.g., trees, types of fruit, and so on). Indeed, such considerations led J. J. Gibson to emphasize the role of surfaces and texture fields in perception rather than the edge-based properties considered by Lowe. My own view is that the visual system is very opportunistic: Depending on the objects one to distinguish, one encodes different kinds of information. A problem with this idea, however, is that one cannot know in advance what will be useful. To distinguish a tiger from a leopard, stripes are the key; but one does not thus look for stripes on every object one sees.

Such considerations have led me to revise my characterization of what the preprocessing subsystem does. I now suspect that it groups edges and regions using two kinds of principles. First, following classical Gestalt theory, the subsystem must use some bottom-up processes to group input, forming groups like those noted by Lowe but also grouping areas of similar color and texture into regions. In the previous theory these functions were carried out in part by a "feature detection" subsystem; I no longer see principled reasons to assume that such a distinct subsystem exists. It is likely that different "channels" exist in the preprocessing subsystem (e.g., see Corbetta, Miezin, Dobmeyer, Shulman, & Petersen, 1990), but the information ultimately is used together to define perceptual units.

Second, I assume that the subsystem can be "tuned" via top-down "training" to organize material. That is, the preprocessing network receives feedback from higher areas so that it can more easily encode visual characteristics that have proven useful in the past. These characteristics can be anything, ranging from a peculiar colored splotch, to a pattern of light intensity, to a configuration of bumps on a surface; an oddly shaped blotch on a cushion may be just the thing to distinguish one's chair from others of the same type.

Biederman and Shiffrar (1987) describe an unusual example of perceptual learning that seems to rely on this sort of opportunistic encoding. They found that subjects could learn to evaluate the sex of day-old chicks once they learned how to attend to the shape (convex versus concave or flat) of a particular cloacal structure. My view is that perceptual learning actually alters the way we organize perceptual input, changing processes in the preprocessing subsystem. Kosslyn (1987) sketches out an algorithm for such perceptual learning, a variant of which was implemented elegantly by Jacobs, Barto, and Jordan (in press).

The preprocessing subsystem would be used in imagery as part of "image inspection," particularly when imaged objects have been combined in novel ways. In this case, perceptual organizations produced by the subsystem would play a critical role in the matching processes that are carried out in a subsequent subsystem as well as in image retention (described below).

Motion relations. Kosslyn et al. (1990) did not consider an important source of information used to identify objects: characteristic patterns of movement. Such information is used in two ways. First, the visual system can infer "structure from motion." Fragments that move in the same way are grouped together. This organizational principle is very powerful (e.g., Ullman, 1979). Second, motion provides characteristic cues that can be used to identify objects. For example, Johansson (1950, 1975) noted that we can recognize a human form solely on the basis of the patterns of movements of its joints, and Cutting and Kozlowski (1977; see also Cutting and Proffitt, 1981) reported that people can recognize individuals solely on the basis of such information. In addition, it has long been known that neurons in some of the higher visual areas of the macaque respond selectively to different patterns of motion. For example, some neurons in the inferior temporal lobe respond selectively to different patterns of gait (e.g., Gross, Desimone, Albright, & Schwartz, 1984).

Because the computation of motion relations is distinct from the kinds of computations necessary to organize static perceptual units, I posit a distinct motion relations subsystem. Whereas the preprocessing subsystem organizes shapes into perceptual units, the motion relations subsystem extracts key aspects of motion fields, and sends this information to a visual memory (to be discussed in the

following section) in which previously encountered motion patterns have been stored. This subsystem is used in imagery in the same way it is used in perception, allowing one to detect previously unnoticed patterns of movement in remembered or novel images.

Pattern activation. Visual memories must be stored somewhere in the system, or recognition could not take place; recognition, by definition, is the matching of input to stored information. Kosslyn et al. (1990) infer a pattern activation subsystem in which visual patterns are stored; these patterns correspond to shapes of objects or parts of objects. We hypothesized, based on results from nonhuman primates, that the pattern activation subsystem is implemented in the inferior temporal lobes.

Each visual memory is composed of a set of perceptual units (positioned in specific locations) and a set of motion relations. The pattern activation subsystem receives both sorts of information as inputs; perceptual units are organized by the preprocessing subsystem and motion relations are extracted by the motion relations subsystem. Both sorts of inputs are matched to the corresponding types of information stored in the visual memory. If both sorts of properties match those associated with a single stored pattern very well, this match is sufficient for object recognition.

Kosslyn et al. (1990) assumed that matching to stored information was performed using the *viewpoint consistency constraint* (Lowe, 1987b). According to this principle, the precise orientation or location of the perceptual units is irrelevant; all that is critical is that the configuration of perceptual units be consistent with seeing an object from a single point of view. This idea fails, however, to account for the wealth of data showing that pictures are more difficult to name in some orientations than others; indeed, the time to name a picture increases with the angular disparity from the upright, with a slight dip in this increase when it is upside down (for a review, see Jolicoeur, in press).

The fact that pictures require more time to name in various orientations suggested to Jolicoeur (in press), Tarr and Pinker (1989), and others that the representation is viewer-centered. Furthermore, they conjecture that viewer-centered input representations are matched directly against these stored representations. One of these two assumptions, about the stored representation or the matching process, must be incorrect, if only because memory for the left-right orientation of pictures is extraordinarily poor (e.g., Nickerson & Adams, 1976). Indeed, when people are asked to name previously seen pictures of objects, they identify mirror-reversed pictures as easily as the originals (Biederman, unpublished data). In fact, Kosslyn and Park (1990) showed that incidental memory for left-right orientation is at chance when previously memorized pictures are subsequently presented to the left visual field/right hemisphere, but are better than chance when they are presented to the right visual field/left hemisphere. This dissociation suggested to us that memory for left-right orientation is accessed separately from the representation of shape per se. (Indeed, if the left dorsal system is in fact better at specifying categorical spatial relations, as discussed below, the result is easily interpreted.)

The sensitivity to planar orientation can be reconciled with the insensitivity to left-right orientation if the stored representation is viewer-centered, but the matching process exploits the viewpoint consistency constraint. In this case, the viewpoint consistency constraint has only limited power, because it is used to match input to a restricted set of information in long-term memory (not a full three-dimensional model, as posited by Lowe). Although use of the viewpoint consistency constraint would match a pattern equally well to itself and its mirror reversal, the sensitivity to planar orientation would result because perceptual inputs are organized differently depending on how a stimulus is oriented in the plane. For example, Rock (1973) provided compelling demonstrations that forms are organized at least in part with reference to their gravitational upright. When an object is oriented oddly, at least some of its components may be organized differently in the preprocessing subsystem--and so will not match the information stored in the pattern activation subsystem. It is of interest that most of the effects of orientation are eliminated if a person is warned that an object may appear at an odd orientation--presumably because subjects over-ride the default gravitational coordinate system and instead organize portions of the object relative to each other (for a similar idea, see Jolicoeur, in press).

In short, I am proposing that Lowe's viewpoint consistency constraint must be understood in the context of the effects of orientation on perceptual organization. Depending on the orientation, a pattern is organized into different units, and subsequent matching is between such units.

The claim that shapes are matched using the viewpoint-consistency constraint seems to contradict properties of neurons in the inferior temporal lobe. For example, Perrett et al. (1984) present

good evidence that many neurons in the inferior temporal lobe not only are selectively tuned for faces, but also respond selectively to faces seen from particular points of view. Some neurons, for example, respond to the left profile of a face but not to the right, and others respond only if the eyes are pointed in a specific direction. My view is that Perrett et al. may be recording from an area that is used to direct action; this area is near the posterior portion of the superior temporal sulcus, which has rich interconnections to the parietal lobe. This portion of the parietal lobe has a role in directing action (Andersen, 1987; Harries & Perrett, in press, appear to adopt a similar perspective). Viewer-centered information clearly is necessary to guide reaching and other movements. There is no evidence, to my knowledge, that these cells are involved in recognition per se.

If an input does not match any representation very well or matches more than one stored representation to the same degree, additional processing is necessary. In this case, Lowe (1987a, b) found it useful to project back an image of the best-matching object, and then to compare this image, template-style, to the pattern in the input array (which corresponded to our visual buffer; see also Ullman, 1989). The image was rotated and its size adjusted until it matched the input as well as possible; this adjustment process may partially account for the increased time to name misoriented objects (Jolicoeur, in press). This operation is interesting in part because it suggests that imagery may have grown out of mechanisms that evolved to match stored representations to inputs during perception, and once it was available it was then used in other contexts.

Images of individual shapes, then, are formed by activating visual memories top-down, and this process in turn induces a pattern of activity in the visual buffer (Kosslyn, 1987). The areas that presumably are involved in storing visual memories are not topographically organized (Van Essen, 1985), and many of the geometric properties of stored shapes may be only implicit (not explicit) in the representation. By analogy, a list of coordinates does not make all information about collinearity explicit, but such information is implicit in the representation. In order to make local geometric relations explicit, it is necessary to use such stored information to produce a representation in an array format.

Furthermore, according to the present formulation, information about motion is implicit in the long-term memories stored in the pattern activation subsystem; to reinterpret motion, these representations must be unpacked in an image. Again, it is the geometric properties of the visual buffer that allow this information to be made explicit and hence subject to new interpretation; motion is registered by systematic shifts of points from location to location in the visual buffer.

The activation of a visual memory is but one component of visual image generation. As noted earlier, we can create composite images, which requires combining stored memories in novel ways. Furthermore, in some cases we mentally "draw" new patterns, "seeing" shapes that do not correspond to individually stored perceptual units. In order to understand these abilities, we need to consider additional components of the system.

Subsystems of the dorsal system

A second major cortical pathway projects dorsally from the occipital lobes, up to the parietal lobes. The usual description of this pathway is that this "dorsal system" is concerned with spatial properties, such as location, size, and orientation (see Maunsell & Newsome, 1987). Indeed, Ungerleider and Mishkin (1982) identify the ventral and dorsal systems as being concerned with "what" and "where," respectively. I infer that the dorsal system receives information from the attention window at the same time as the ventral system; hence, both systems are computing information about the contents of the same region of the visual buffer.

I have recently revised my thinking about the role of the dorsal system, in large part on the basis of findings in nonhuman primates. As Andersen (1987) and Hyvarinen (1982) point out, a pervasive property of neurons in the posterior parietal lobes is that they fire prior to the animal's initiating a movement or are sensitive to the consequences of a movement. The parietal lobe appears to be concerned in large part with controlling and monitoring movement, and spatial information must be encoded to serve these ends.

The idea that the parietal lobes are not simply concerned with encoding spatial properties, but rather with encoding information to guide action, may help to clarify a longstanding puzzle: In the experiments by Pohl (1972) and Ungerleider and Mishkin (1982), monkeys discriminated between patterns on food lids or between the locations of a small "landmark." When the animals' parietal lobes

were removed, their performance on the landmark task was devastated, but they performed the pattern task well; this result is consistent with the idea that the parietal lobes are critically involved in encoding location. In contrast, when animal's temporal lobes were removed, their performance on the pattern discrimination task was devastated, but they performed the location task well; this result has been taken to show that the temporal lobes encode shape.

A problem with these interpretations is that spatial properties of the patterns in the shape task are often sufficient to discriminate among them. For example, a monkey may have had to discriminate between checks and stripes; in this case, there were fewer locations defined by the stripes than the checks, the patterns had different sizes, and they had different orientations (Holmes & Gross, 1984, showed that animals can discriminate orientation even when the temporal lobes are removed). Thus, all of the spatial properties of the patterns were sufficient to discriminate between the patterns. And yet monkeys without temporal lobes are severely impaired at the discrimination—even when the parietal lobes are intact.

I have puzzled over this apparent paradox for years, and only recently had a hint of a possible resolution from the behavior of a patient studied by Kosslyn, Daly, McPeck, Alpert, and Caviness (1990). As is summarized in the second part of this report, this patient had suffered damage to the left frontal lobe and had hypometabolism (revealed by PET scanning) in the occipital-temporal area on the left side. We asked this patient to discriminate between patterns that were formed by filling in cells of a 4 x 5 grid. He had some difficulty encoding patterns, and reported that he remembered the patterns in grids by *looking at each individual cell*. He apparently remembered the patterns as sets of filled locations in the grid. And in fact, the more segments the pattern had, the more time he required to encode them. When the grid lines were removed, so that cells were not clearly defined, he could not use this strategy and his response times changed accordingly; there now was no effect of the number of segments on the time to encode patterns. This difference in response times suggests that the patient was not making the same pattern of eye movements when viewing both types of displays.

One way to understand these results is to infer that the location information is normally encoded in a form suitable for directing action, and can only be used for recognition by making eye movements and recoding the location information into a different format. Think about how easy it is to toss an object into a wastepaper basket, compared to how difficult it is to estimate the distance of the basket from you. I have informally tested a series of people who enter my office, and found that some can throw better than they can estimate the distance and vice versa for others. The important claim is that there is a dissociation between the two kinds of information. This observation makes sense if the information about location is "encapsulated," and can only be directly used to guide action. McLeod, McLaughlin, and Nimmo-Smith (1985) provide good evidence for such a dissociation.

If so, then the monkeys without temporal lobes may have been unable to discriminate between patterns because they did not hit on the strategy of moving their eyes over the patterns, which would have allowed them to encode the spatial properties in a way useful for identification. It would be interesting to observe whether monkeys without temporal lobes could discriminate between checks and stripes if they had been trained to look at the dark regions of patterns prior to surgery.

Kosslyn et al. (1990) did not consider the idea that the parietal lobes encode spatial information in a format to be used to guide action. This idea leads me to modify Kosslyn et al.'s characterizations of the subsystems in the dorsal system.

Spatiotopic mapping. Location information is specified relative to the retina in the visual buffer (these maps are retinotopic; see Van Essen, 1985). Because a retinotopic representation changes whenever one moves one's eyes, it is not useful for object identification, navigation, or tracking. One needs a representation of an object's location relative to another object or part, not relative to the retina. Andersen, Essick, and Seigel (1985) found cells in area 7a (part of the parietal lobe) of the macaque that respond to location on the retina, as gated by eye position, and Zipser and Andersen (1988) showed that the outputs from sets of these neurons are sufficient to indicate location relative to the head.

I therefore infer a subsystem that receives as input a retinotopic position and the positions of the body, head, and eyes, and computes where an object or part is located relative to other objects or parts. During both vision and imagery, the output from the spatiotopic mapping subsystem is a set of spatiotopic coordinates that are tailored to guide action. Kosslyn et al. (1990) assumed that these coordinates were general purpose representations, but the present view is that they are dedicated for

use in guiding actions. This idea has implications not only for how we form images, but for how we decode spatial information from images, as noted below.

Coordinate spatial relations encoding. We often want to store spatial information in memory. For example, to navigate efficiently in familiar rooms, it is useful to store the locations of furniture. This can even allow one to navigate in the dark. Thus, I hypothesize the existence of a subsystem that encodes the types of coordinates used to guide action. This subsystem does not encode motor programs, but rather coordinates that can be used to guide actions.

Fisk and Goodale (1988; see also Goodale, 1988) found that right-hemisphere damaged patients had difficulty in initiating a movement when asked to point at a dot. This result is consistent with the idea that the right hemisphere has a special role in encoding the coordinates that are used to guide actions. A key component of such computation is the precise specification of the location of an object, and hence it is of interest that Hellige and Michimata (1990), Koenig et al. (1990), Koenig, Reiss, and Kosslyn (1990), and Kosslyn, Koenig, Barrett, Cave, Tang, and Gabrieli (1989) provide evidence that the right hemisphere can encode metric spatial information more effectively than the left (see also De Renzi, 1982).

This subsystem can be used in imagery in at least two distinct ways: It can play a role both in image generation when multiple parts are assembled, and in image inspection, encoding spatial relations among parts of imaged objects; these roles will be discussed shortly.

Categorical spatial relations encoding. Different tasks require the use of different types of spatial relations. Consider the situation in which one is so close to an object that one only sees a small portion of it in a single fixation. In this case, the ventral system would identify parts, and the spatial relations would be encoded via the dorsal system. Many objects, such as a human form, can assume a wide range of positions as the parts move. In order to identify such objects, the spatial relations among the parts should be specified rather abstractly. The fact that the forearm is "connected to" the upper arm remains true no matter how the metric relations between them vary.

The *categorical* spatial relations encoding subsystem encodes relations such as "connected to," "left of," "under," or "above." These representations capture what is stable across instances that may differ in terms of precise metric relationships. As Kosslyn, Chabris, Marsolek and Koenig (in press) review, previous work provides evidence that this subsystem is relatively more effective in the left cerebral hemisphere. This finding is consistent with the long-standing reports that Gerstmann's syndrome, which includes left-right confusion as one component, occurs following damage to the left angular gyrus (e.g., see DeRenzi, 1982).

A reinterpretation of the distinction. Sergeant (in press) reports that the hemispheric dissociation between coordinate and categorical encoding only occurs when the stimuli are displayed at relatively low contrast. This result puts real pressure on the theory of Kosslyn (1987) and Kosslyn et al. (1990), and has caused me to reconceptualize the theory. The driving force behind the revised conception is a recent finding by Kosslyn, Hillger, Livingstone, and Hamilton (1990).

We asked subjects to view two short line segments presented in succession and to decide whether the lines had the same orientation. Both segments were presented in the same visual field while the subject stared at a central fixation point. The important variable was the distance between the locations of the lines in each pair; they were either relatively close (within 1° of visual angle) or far (up to 8° apart). When the segments were relatively close together, subjects were more accurate if the stimuli were presented initially to the left hemisphere; when they were relatively far, subjects were more accurate if the stimuli were presented initially to the right hemisphere.

One account of these findings hinges on the idea that neurons in the high-level visual areas in the two hemispheres have different sized receptive fields, perhaps because they receive input from different retinal ganglia. It is possible that some of the ganglia, such as the magnocellular neurons (see Livingstone & Hubel, 1987), have a special role in "preattentive" processing. The magnocellular ganglia encode motion and flicker very well, which is useful for guiding eye movements and subsequent "focal" attention (see Neisser, 1967). Furthermore, the magnocellular ganglia have relatively large receptive fields, which would help preattentive processing to monitor the entire visual field.

The finding that the right hemisphere encodes spatial location better than the left follows directly from the idea that the right hemisphere monitors larger, more overlapping receptive fields:

Computer simulation modeling has shown that relatively large overlapping receptive fields are more effective at using "coarse coding" to register the location of a dot relative to a line than smaller, less overlapping receptive fields (Kosslyn, Chabris, Marsolek, & Koenig, in press). This notion appears to be consistent with Sergent's own interpretation of her results. In contrast, our computer simulations showed that smaller, less overlapping fields are more effective for dividing space into discrete bins, which correspond to some spatial relations categories (such as above/below or left/right). This idea, then, leads us to expect a *left-hemisphere advantage only for some categorical spatial relations*, namely those that allow space to be carved into discrete regions. Preliminary results in our laboratory suggest that this prediction is worth taking seriously.

The idea that the left hemisphere typically monitors smaller local regions than the right hemisphere is consistent with numerous findings. For example, the left hemisphere plays a critical role in encoding portions of objects, whereas the right hemisphere plays a critical role in encoding global patterns (e.g., Delis, Robertson, & Efron, 1986). Furthermore, people categorize parts of objects faster when the objects are shown initially to the left hemisphere, whereas they categorize overall shapes faster when they are presented initially to the right hemisphere (see Van Kleeck, 1989, for a review). Although large overlapping receptive fields are good for encoding location, they are not as good for encoding shape. For this purpose, smaller receptive fields provide greater resolution (because they average input over smaller areas).

Thus, the revised theory leads us to expect differences in the ventral systems in the two cerebral hemispheres. Kosslyn (1987) alluded to such possible differences, but did not provide detailed arguments for them. Specifically, the notion that the higher visual areas of the two hemispheres differ in the sizes of the receptive fields they monitor implies that the contents of the pattern activation subsystem may also differ: The left hemisphere may store better representations of separate portions of objects, whereas the right may store better representations of overall shapes.

The claim that the left hemisphere encodes portions of objects more effectively than the right might help to explain another of Fisk and Goodale's (1988) findings: Patients with left hemisphere damage could initiate a *reaching movement normally*, but had trouble controlling it (particularly in the deceleration phase). Reaching apparently has two phases, initiation (which is open-loop) and fine-tuning (which uses feedback). The right hemisphere may be critical in the first phase because it computes the location of the target better. And if the left hemisphere is more adept at encoding portions of objects, it may be critically involved in orchestrating the second phase of a reach; we typically reach for a portion of an object, such as the handle of a cup or the bottom segment of a pen.

Now let us return to Sergent's finding that the right-hemisphere advantage for encoding spatial coordinates depends on the level of contrast. Our computer simulations showed that if high contrast allows more input units to fire, the differences in the sizes of receptive fields no longer effect the ease of computing either metric distance or discrete bins. When very many units contribute, many of them have overlapping receptive fields and many do not. Thus, the networks can map both functions easily.

To summarize, the revised theory of categorical versus coordinate spatial relations encoding rests on the idea that the right hemisphere monitors larger receptive fields than the left, which is useful for detecting stimuli over the entire field. This information in turn is used to direct movement (such as head and eye movements towards a stimulus). These large fields overlap, conferring high resolution for specifying position via coarse coding. In contrast, by monitoring smaller receptive fields, the left hemisphere is better able to focus in on important characteristics of an object. These smaller receptive fields are also useful for carving space into bins, which may correspond to some types of categorical spatial relations. The differences between the hemispheres are a matter of degree, and when contrast is very high large amounts of all types of input are sent to both hemispheres, minimizing the differences. [Footnote 2]

Like the coordinate spatial relations encoding subsystem, the categorical spatial relations encoding subsystem can be used in imagery in at least two distinct ways, as described below.

Associative memory

The simple fact that people can report from memory where furniture is placed in their living rooms indicates that the outputs from the dorsal and ventral systems are conjoined downstream. Kosslyn et al. infer an associative memory in which such conjunctions are stored. If an object is seen close

up, so that it is examined over the course of multiple eye fixations, then associative memory will be used to build up a composite representation of the object and to identify it. During perception, the outputs from the ventral and dorsal systems are matched in parallel in associative memory to parts and relations of stored objects. The system converges on the identity of the object being viewed by finding the stored representation that is most consistent with the encoded parts and their spatial relations. When such evidence exceeds a threshold (which presumably can be varied, depending on context), identification occurs.

Goldman-Rakic (1987) summarizes evidence that one aspect of associative memory involves structures in the frontal lobes. In particular, she shows that area 46 in the dorsolateral prefrontal lobes is critically involved in storing memory for location. If this area is damaged in one hemisphere, an animal cannot retain in short-term memory the locations of stimuli in the contralateral field. The area is topographically organized; when different portions are damaged, memory is subsequently impaired for different regions of the visual field. Furthermore, Goldman-Rakic shows that areas of the parietal lobes that are involved in encoding spatial properties not only project to the frontal lobes, but also receive rich projections from them.

Associative memory plays a critical role in imagery for at least two reasons. First, this is where information is associated with an object's name. We often form images upon hearing the name of objects. Second, because associative memory integrates the outputs from the dorsal and ventral systems, it must contain representations of the structure of scenes and objects. To image an object that is composed of more than one part, we must access information about the structure of the object and use this information to activate the appropriate visual memories and the appropriate spatial relations representations. This process involves additional subsystems, as noted below.

Subsystems used in top-down hypothesis-testing

We see only about 2° of visual angle with high resolution. Thus, we often must move our eyes over an object or scene during recognition and identification. Logically, there are only three ways in which we can guide eye movements: randomly, on the basis of bottom-up information (e.g., motion), or using stored information. Yarbus (1967) provides ample evidence that knowledge is often used to guide one's sequence of attention fixations. Kosslyn et al. (1990) inferred a set of subsystems that are involved in accessing and using stored information to shift attention.

Coordinate property lookup. Often, the location of objects in a scene or the locations of parts on an object are important in identification. Thus, Kosslyn et al. (1990) postulate subsystems that can access stored information about the spatial arrangement of parts of objects and can use this information to shift attention to relevant locations. The present revision of the theory leads me to characterize the coordinate property lookup subsystem slightly differently from Kosslyn et al.; it accesses stored information that can be used to guide movements precisely. A subsystem that accesses such stored information appears to be implemented in the frontal lobes, near the frontal eye fields (area 8; cf. Luria, 1980).

The coordinate property lookup subsystem seems to be involved in many image generation tasks. For example, if asked to describe where the furniture is in their living rooms, most people move their eyes and report scanning to a location in an image and "seeing" the object. One interpretation of this finding is that the furniture is in fact not present until one scans to the appropriate location, and that such scanning involves activating motor-based coordinate representations of location. These representations are useful for guiding action, and in order to recover a representation of a specific location one must activate a motor program. One often may be able to inhibit the actual execution of the program, but perhaps not completely. Hence, one often moves one's eyes in the course of building up the image (cf. Hebb, 1949).

Categorical property lookup. Categorical representations group positions and treat them as equivalent; in contrast, coordinate representations specify the finest possible distinctions. Hence, the two representations are qualitatively distinct, and Kosslyn et al. (1990) argue that they logically require different operations to access. Thus, Kosslyn et al. (1990) infer a second lookup subsystem that accesses stored information about the categorical locations of objects in a scene or individual parts. This subsystem may also be implemented in the frontal lobes. [Footnote 3]

This idea implies that there are two distinct ways of adding parts to an image, one using coordinate spatial representations and one using categorical spatial representations to specify the

parts' locations. If one images one's living room repeatedly, I have observed, one no longer moves one's eyes. It is possible that with repeated use, the motor-based coordinate representation is recoded into a categorical representation. Indeed, Koenig, Kosslyn, and Chabris, and Gabrieli (1990) found that the right-hemisphere superiority for metric judgments disappears after practice, which is consistent with this idea.

Attention shifting. Recent evidence suggests that the human visual system probably includes at least three subsystems that are used to shift attention: One that disengages attention from the current location (which appears to involve the parietal lobes); one that shifts attention to a new location in space (which appears to involve the superior colliculus); and one that engages attention at that new location (which appears to involve the thalamus; see Posner et al., 1987). Kosslyn et al. (1990) chose a coarser level of modeling in which all attentional control mechanisms were grouped into a single attention shifting subsystem.

The attention shifting subsystems guide the movement of the body, head and eyes, and also adjust the attention window in the visual buffer (both in perception and visual mental imagery). These mechanisms are important for several reasons. First, they guide image scanning and zooming. Second, they play a critical role in some forms of image generation. Consider, for example, a task developed by Podgorny and Shepard (1978). They showed people empty 5 x 5 grids, and asked them whether a dot or dots would be covered if a specific block letter were present in the grids (the subjects saw the block letters in advance, which were formed by selectively filling in cells in the grid). In this task, one selects specific cells to pay attention to; one does not activate stored visual memories.

The idea that images can be formed by allocating attention also allows us to consider "mental drawing." One can image a line simply by shifting attention over the visual buffer and activating each small region of the buffer in turn. This process will create a representation of a "path" in the visual buffer, which in turn can be processed just like any other pattern of activity (such as those arising during perception).

Thus, we are led to make another new distinction: Some forms of imagery involve activating *stored visual memories*, whereas others involve engaging attention in specific regions. This distinction leads to a simple prediction: There is no reason why the complexity of an object need affect the time to image it using the first method; if the object is stored as a single perceptual unit, the unit is simply activated. For example, a normal face might be easier to image than a face with scrambled features, even though both have the same number of features. The normal face has been seen so often that there may be perceptual grouping processes built into the preprocessing subsystem that produce a single representation of a face, which can be imaged as such; in contrast, the scrambled display cannot be encoded as a single unit, and hence multiple units are encoded and must later be imaged individually. The other sort of imagery does not offer this possible difference; because the attention window can only pick out a regular region in the visual buffer, one will always need to shift it to attend to different regions—and so more time always will be required to image patterns that contain more component parts.

Thus, when imaging a letter in a grid, for example, one will need to attend to each segment in sequence. The more segments in the letter, the longer it should take to form the image. Kosslyn, Cave, Provost, and Von Gierke (1988) confirmed this prediction. In contrast, if one's eyes are closed and one is merely imaging what a previously seen letter looks like, there is no reason to expect that more segments should result in longer times; one simply activates the visual memory. Kosslyn, Hillger, Engel, Clegg, and Hamilton (1990) have confirmed this prediction.

Transformation

Lowe (1987a, b) proposed that when nonaccidental properties do not match the input very well during perception, an image is generated and matched to the input pattern. Lowe's computer vision system tried to maximize the match by rotating the generated image and adjusting its size scale. I have adopted his use of imagery in object recognition, which leads me to predict that there should be two distinct ways of imaging movement. First, if one has stored a visual memory of a moving object in the pattern activation subsystem, it can simply be reactivated. For example, imaging a horse running is simple if the visual memory itself contains information about its movement patterns. This information is purely visual.

Second, if the object was encoded without motion information, this information can be added by changing the spatial representations encoded in the dorsal system. In many cases, the only available

representation of location, orientation, and size is encoded in a motor format. In these situations, one must execute motor operations on these representations to alter them. This idea predicts that people sometimes will perform implicit motor movements when transforming shape.

The two kinds of motion information may often be used together. One may not have encoded a pattern of movement over a length of time, but instead registered a succession of moving images. In this case, one will move one's eyes when replaying the image, with the eye movements indicating that the relative locations of the separately encoded images has been activated in the course of integrating them.

When one transforms an object that was not seen moving, one must actually alter the image in the visual buffer. When a three-dimensional object is rotated, new portions of the object will come into view and hence new visual memories must be activated. Thus, it is of interest that there are rich connections between area 7a in the parietal lobe and the regions of the inferior temporal lobe that presumably underlie visual memory (Harries & Perrett, in press). As one changes the spatial properties of the object, this in turn alters the aspects of the visual memory that are projected back to the visual buffer.

Summary and Critical Distinctions

The logic used to develop the theory of imagery hinges on the idea that perceptual mechanisms are used in imagery. Thus, I will summarize the way the system operates during perception proper before showing how it can provide accounts for the five key imagery phenomena reviewed at the outset.

Identifying objects

An object is identified by first positioning the attention window in the appropriate part of the visual buffer. Once the image of the object is enveloped by the attention window, it is sent simultaneously to the dorsal and ventral systems for further processing. The ventral system, which encodes object properties, attempts to organize perceptual units and match them to those of stored shapes. The dorsal system, which encodes spatial properties, converts retinal location to spatiotopic coordinates and encodes categorical spatial relations and motor coordinates. An object can be recognized at first glance if the match to a stored shape in the ventral system is very good. However, if the match does not definitively implicate a single object, then the identity of the closest matching object is treated as an hypothesis to be tested.

Hypothesis testing is done by accessing properties (such as parts or distinctive marking) and spatial relations between the properties of the candidate object stored in associative memory, and then positioning the attention window at the location of a sought property. The portion of the image at that location is then encoded via the ventral and dorsal systems. The subsequent output of these systems, which is sent to associative memory, may provide evidence in favor of the hypothesis or may lead to the formulation of a new hypothesis. The top-down hypothesis-testing cycle is repeated as many times as necessary until the stimulus has been identified (see Kosslyn et al., 1990, for details and computer simulations).

Imaging objects

The imagery phenomena considered earlier are explained in the following ways.

Geometric representation. The visual buffer functions to make explicit the local geometry of surfaces of objects. An image is a pattern of activation in topographically organized areas, and so portions of the representation correspond to portions of the object.

Generation. Images of single remembered shapes (that may or may not include color, texture, or motion characteristics) are formed by activating stored visual memories in the pattern activation subsystem; this process results in a pattern of activation in the visual buffer, which is an image representation. In addition, we are led to posit four distinct types of image generation that are used when multiple parts are amalgamated or novel patterns are formed, defined by a two-by-two table: One either activates visual memories or allocates attention, and positions portions of the pattern using either categorical or coordinate representations of spatial relations. Consider first image generation when one activates visual memories of shapes, as occurs if one images a familiar scene. In this case, a description of the scene would be accessed in associative memory. This description would specify the objects and their spatial relations. Each object representation would in turn be used to activate a visual memory in the pattern activation subsystem, and the appropriate spatial relations representation would be used to position it correctly. If a coordinate spatial relations representation is encoded, a

motor program is activated and the results used to compute the location; I assume that the categorical spatial relations encoding subsystem then is used to encode the spatial relation into associative memory, where it is then used to position the image appropriately. The process of positioning the component object involves shifting the attention window to the appropriate region of the visual buffer, and forming the image at that location. (Recall that I assume, following Lowe, that the process of forming images can be adjusted to produce them in different locations in the visual buffer.) If a categorical spatial relations representation is stored, it can be used immediately to position the attention window and then form the image of the object or part in that location. [Footnote 4]

This sort of image generation may result in increased amounts of time for more complex objects. We expect such increases if objects are stored as separate visual memories of their constituent parts and each spatial relation specifies a part's location relative to a different part; hence the other part must be present before the new part's location can be computed. In principle, there is no reason why multiple parts cannot be imaged at the same time if their locations are specified relative to the body or another independent reference point. In each case, the imaged patterns may be static or moving.

The second sort of image generation is similar, except that no stored visual memories of patterns are activated. In this case, one simply picks out portions of the visual buffer to be activated. This process is done by guiding the attention window to different regions, which can be accomplished using either categorical or coordinate stored spatial relations representations.

Inspection. Objects in images are "inspected" using the exact same mechanisms as in perception. The pattern of activation in the visual buffer is surrounded by the attention window, and information is sent to the ventral and dorsal systems, as described above. These processes allow one to examine previously unconsidered shapes, colors, and textures as well as locations, orientations, and sizes. In addition, patterns of motion in the image can be encoded using the motion relations subsystem.

Recoding. The same processes are used in perception and imagery to store a new pattern in the pattern activation subsystem (i.e., enter a new visual memory) or in associative memory (i.e., enter a new structural description). I do not have a theory of how these processes operate, but the fact that the same subsystems and representations are used in the two types of processing implies that whatever mechanisms are responsible for learning in perception will also allow learning in imagery.

Maintenance. Image maintenance can be considered as a special case of image generation, with the generation mechanisms simply being used repeatedly to refresh an existing pattern of activation in the visual buffer. If a novel pattern is created, one must first encode the pattern into the pattern activation subsystem, and then activate this new representation to recreate the image. To the extent that one can "tune" the preprocessing subsystem to organize information into fewer units ("chunks") before these visual memories are created, one can hold more information in a single image.

The process of image maintenance plays a critical role in one form of "working memory" (Baddeley, 1976). In my view, there are three types of memory in the system: Short-term memory is the use of a perceptual buffer to represent information activated from long-term memory. The visual buffer is an example of such a short-term memory. Long-term memories may also be modality specific or may be amodal (i.e., in associative memory). The pattern activation subsystem is an example of a modality-specific long-term memory. Working memory is a) the combination of the information being held in the various short-term memory structures and the information that is activated in the various long-term memory structures, and b) the "control processes" that activate information in long-term memory and allow information to decay in short-term memory. That is, there is a dynamic relation between short-term and long-term memory. More information typically is activated in long-term memory than can be represented in short-term memory, and hence there often is a complex "swapping" process between the two types of structures at work, shuffling information in and out of short-term memory. Presumably the frontal lobes play a critical role in governing this swapping process, just as they do in selecting objects to be imaged. Note, however, that "loading up working memory" may consist of loading up the short-term buffers, which would not necessarily influence information stored in long-term memory.

Transformation. Finally, the revised theory leads us to expect that there are two distinct ways of transforming imaged patterns. First, if motion was an intrinsic part of a visual encoding, it can be recreated simply by activating the visual memory. Second, the spatial relations representations in the dorsal system can be altered, in part by running motor programs. This kind of operation is very flexible, and can be applied to a wide range of objects.

II. Using the Theory to Diagnose a Deficit

We have used the theory to guide research of a variety of types, ranging from divided-visual-field studies of normal subjects to studies of focal lesion patients. The latter sort of work is arguably the most innovative, and thus I will focus on it here (in the course of summarizing the theory I described several typical divided-visual-field studies we performed with the support of the grant). We use chronometric techniques developed in cognitive science to delineate the pattern of deficits in a single patient; this was particularly interesting because this patient had a focal lesion in the frontal lobe, which disrupted connections in the region of the sylvian fissure; some of the disconnected areas are thought to be involved in vision, and hence we expected our patient to have visual deficits.

Logic of the experiments

We began by documenting that the patient did indeed have a visual deficit, and then conducted a series of 16 experiments to discover how the system had been disrupted. Each experiment was designed so that normal performance can be achieved only if a particular subsystem operates normally. We were able to implicate an individual subsystem by observing relative performance in two conditions. That is, any given task draws on multiple subsystems, and hence the overall performance score for a given task reflects the operation of numerous subsystems. To focus on particular subsystem, we identified a variable that should affect only the operation of that subsystem. We then manipulated that variable to force the subsystem to engage in more processing, and observed the consequences on performance. If the subsystem is normal, then forcing it to work harder should produce decreased performance comparable to that found in normal subjects. If the subsystem is impaired, however, then forcing it to work harder should produce marked dysfunction.

We measured performance by examining the relative differences in response times and error rates, comparing a relatively "easy" and "difficult" version of each task. It is important to realize that response time and error rate are inter-related. If a subsystem is impaired, a subject could try to respond in a normal amount of time and hence would produce many errors. Or, a subject could engage in more thorough processing, in which case he or she might not produce many errors, but would take much longer to respond. This "speed-accuracy trade off" function has been well studied in cognitive science (e.g., Luce, 1986).

The logic underlying our task design has been used by researchers who study mental rotation in brain-damaged patients by examining the slopes of mental rotation functions (e.g., Kosslyn, Berndt & Doyle, 1985); this research rests on the assumption that when a stimulus is presented at a greater angular disparity, a mental rotation process must perform additional processing to reorient the representation. Because all other aspects of the task are held constant, the effects of manipulating angular disparity (i.e., the slope of the function relating angle and response time) can be taken to reflect the efficiency of the mental rotation process.

Unfortunately, although this logic can allow us to eliminate alternative hypotheses, it cannot directly implicate an hypothesis: When we find abnormal performance, there almost always will be more than a single possible way to account for it. Thus, we must perform a series of experiments in which we attempt to rule out various possible hypotheses, and then use the pattern of results to interpret instances of impaired performance. In the experiments we have conducted to test our patient, we manipulate variables that can be identified with the operation of a particular subsystem and examine the effects of these manipulations on the patient's scores. In describing each experiment, we begin by outlining the task and then describe the manipulation used to tap into the subsystem of interest.

Patient

The patient, R.V., was a right-handed, bilingual male who had worked in technical training at a large computer company. He had earned a bachelor's degree and was working toward a Master's degree. Six months later he presented with mild anomia and slight deficiencies in speech production. Caplan (1990) tested him extensively on the Caplan-Bub Aphasia Battery, and found that he had moderate comprehension difficulties as well. In addition, he failed to name 13% of simple line drawings of common objects in their picture naming task; virtually all of these objects were animals. He was 39 years old at the time of testing.

Structural Imaging

A CT scan revealed that R.V.'s lesion was focused in the left frontal lobe. The damaged area appeared to be a cone whose base rested on the head of the caudate nucleus and whose tip just touched

cortex near the region of the third convolution of the frontal lobe. Magnetic resonance imaging (MRI) allowed much greater precision in characterizing this focal lesion. The lesion included zones of frank cavitation as well as zones of T1 and T2 signal prolongation consistent with gliosis; it was centered in the centro-sylvian region of the left hemisphere, as is illustrated in the top portion of Figure 2. Its extent was maximum in the region of the frontal operculum where the zone of signal change and cavitation spanned the full thickness of the cerebral wall. At the cerebral surface, the lesion destroyed much of the inferior opercular sections of Brodmann (1909; Bailey and von Bonin, 1951) areas 46, 45, 6, 4, plus the superior extent of 43 within the insula under the rostral parietal operculum. It intruded minimally into fields 3, 1, 2, and 40 within the sylvian fissure. Subcortically, the entire caudate and lenticular nuclei rostral to the thalamus as well as the adjacent segment of the horizontal limb of the diagonal band of Broca's area were destroyed and replaced by cavitation. The intervening corona radiata, external sagittal stium and anterior limb of the internal capsule were either marked by signal change or also frankly cavitated. Involvement of these central white matter systems extended forward through the forceps major beyond the callosal commissure. Caudally, the lesion also destroyed much of the posterior limb of the internal capsule to the level of the pulvinar.

In its extent across the external sagittal stium, the lesion may be assumed to have damaged or destroyed the superior and inferior longitudinal fasciculi as well as the uncinate fasciculus carrying axonal systems linking pre-and postcentral as well as temporal and frontal ipsilateral cortical regions, respectively (Krieg, 1973). In its extent through the coronal radiata, more local ipsilateral cortical interconnections would have been damaged. An estimate of the extent of ipsilateral cortico-cortical denervation, derived by homology from hodological studies in the rhesus monkey (see Pandya & Yeterian, 1985) is provided in Figure 2. Homotopic interconnections of this full region with the opposite hemisphere as well as the connections of cingulate and much of the frontal and orbital cortical regions with the anterior and ventral thalamic nuclear groups and the medial dorsal thalamic nucleus may also be assumed to have been largely interrupted.

Insert Figure 2 About Here

Metabolic Imaging

The fact that R.V.'s lesion apparently disrupted the inferior longitudinal fasciculus raised the intriguing possibility that posterior regions of the brain innervated by this fasciculus might be dysfunctional. To explore this possibility, R.V. was studied with positron emission tomography (PET) to determine local cerebral blood flow and oxygen metabolism.

The PET study was conducted approximately three months after ictus. The scans were performed according to the ^{15}O steady state method (Fackowiack, 1980; Senda, 1988) with a Scanditronix PC-384 positron tomograph (Litton, 1984). R.V. was asked simply to rest with his eyes open while being scanned, performing no particular task. The PET data were transformed to a standardized stereotactic coordinate system using anatomic reference data obtained from XCT according to the method of Talairach (1967). Brain regions were then imposed on the PET data from a digitized version of Talairach's standard stereotactic brain atlas.

The PET scans revealed spatially matched disruptions in flow and metabolism with severe (nearly absent) hypoperfusion and hypometabolism affecting parts of areas 6, 8, 9, and 10, much of areas 44, 45, and 46, the insula and superior aspects of the caudate nucleus. Milder hypoperfusion and hypometabolism were found in the superior temporal gyrus (a portion of area 22), a remote region innervated by the inferior longitudinal fasciculus. Area 8 is clearly involved in shifting eye movements, and area 46 may correspond to a visual short-term memory for spatial location (Goldman-Rakic, 1987). In addition, area 22 is in prestriate cortex, and presumably plays a role in visual encoding (cf. Luria, 1980; Van Essen, 1985).

As is evident in Figure 2, there was a striking correspondence between regions of reduced metabolism as revealed by PET and the regions that are anatomically connected to the damaged location, as revealed by MRI. Thus, we have good reason to hypothesize that some visual-spatial functions should be impaired.

General Method

Control subjects

In order to establish a behavioral deficit, we compared R.V.'s scores on a task with those of a group of 8 control subjects. These subjects were right-handed men who responded to advertisements posted in various locations around Harvard University. They were approximately R.V.'s age (average age 36.6 years, range 33 to 42 years), and each was either working towards a bachelor's degree or had no more than a Master's degree. All subjects, including R.V., were paid for their time. Comparing R.V.'s scores to those from relatively few control subjects will produce conservative estimates of R.V.'s deficits, which is reasonable given the large number of experiments that we must conduct to converge on possible accounts for deficits.

General materials and procedure

All experiments were administered using a Macintosh Plus computer. A Polaroid CP-50 filter was placed over the computer screen to reduce glare, and a chin rest was used so that subjects viewed the displays at a constant distance of 50 cm. Unless otherwise noted, the stimuli (either grids or brackets, as will be described) subtended 3.0° of visual angle horizontally and 3.6° of visual angle vertically, and were centered on the screen.

All subjects were tested in the same conditions in a quiet room with indirect artificial lighting. For every experiment, the instructions, practice trials, and test trials (in that order) were displayed on the screen, and the experimenter read aloud the instructions. The B and the N keys of the keyboard were assigned as the "yes" and "no" response keys, respectively, for all subjects (in addition, for two of the experiments the labels were augmented, as noted below), and the subjects responded by pressing the appropriate key on the keyboard. With patient populations in mind, the response keys were adjacent keys on the keyboard so that all responses were made with just one hand. In all experiments, the subjects were asked to respond as quickly and accurately as possible. The computer recorded both the key pressed and the time taken to make the response; an internal clock was started when a probe stimulus appeared, and stopped when either of two response keys was pressed.

Each experiment began with a practice session, in which all conditions of the experiment were represented at least once. Unless otherwise noted, the practice session consisted of 12 trials that were balanced in the same way as the test trials. The stimuli used in the practice trials were very similar to those used in the test trials, but were not included in the subsequent test trials. During practice trials, the computer beeped and the experimenter repeated the instructions if the subjects pressed the incorrect response key. During the test trials, however, there was no feedback and the experimenter remained silent and out of sight.

In all experiments, no more than three "yes" or three "no" trials appeared in a row. Furthermore, when probe marks were used they appeared equally often on the left and right side of the stimulus. In some tasks, alterations of an initial stimulus were made to produce "no" trials; these alterations also appeared equally often on the left and right sides of the stimulus.

R.V. was tested in three separate sessions by the same experimenter. The first session was 8 March 1989, and the last was 7 September 1989. During every session, R.V. was periodically reminded that he could take a break between experiments; however, he usually choose to continue without breaks. R.V. was easily able to press the prompting and response keys on his own. In addition, he had very little difficulty understanding the task instructions; any difficulties always were quickly sorted out during the practice session at the beginning of the experiment.

When we completed testing R.V., each control subject was individually tested during a single three-hour session; the experiments were conducted in the same order for R.V. and the control subjects. The control subjects were periodically invited to take breaks between experiments. Like R.V., they too usually choose to continue without breaks. (They averaged two 5 minute breaks in the whole three hour session).

The first 8 experiments were conducted in an order designed solely to provide variety and keep the subjects interested. After these experiments were conducted, however, additional ones were designed to pursue specific hypotheses. The experiments were administered in the following order: Ventral Shape Comparison, Shape Comparison, Preprocessing Overload, Short-Term Memory Control, Location Top-Down Search, Scanning, Mental Rotation: Simultaneous Presentation, Location Associative Memory, Preprocessing Followup, Scope of Attention Window, Coordinate Spatial

Relations Encoding, Categorical Spatial Relations Encoding, Pattern Activation Storage, Shape Top-Down Search, Mental Rotation: Sequential Presentation, Pattern Activation Memory.

For ease of exposition, we will present the experiments and results in an order that is logically structured around Kosslyn et al.'s theory of high-level vision. The pattern of results is what is important, and the theory guides its interpretation.

Experiment 1: Shape Comparison

This experiment was designed to document that R.V. had a specifically visual deficit; the task was designed to tap as many of the subsystems as possible, and so if he had a deficit, he should perform abnormally on this task. We knew from earlier testing that R.V. had failed to name 13% of pictures of common objects, but did not know whether this deficit reflected problems in visual recognition and identification per se or problems in accessing or producing names. Thus, to document a visual deficit we designed experiments that were as devoid of semantic content as possible.

In this experiment, the subjects saw a 4 x 5 grid with some of the cells blackened; cells were blackened to form either 1, 2, or 3 perceptual units. The subjects studied the pattern until they had memorized it, and then pressed the space bar. The pattern was removed and there was a brief delay, at which point another pattern was presented. On half the trials, the second pattern was identical to the first, and on half it was modified. The subjects had only to indicate whether the second pattern was the same or different from the first. The *manipulation* here was the number of perceptual units. By varying the number of perceptual units, we taxed the subsystems that encode and store the first pattern and that encode the second and compare it to the representation of the first. The *score* here was the increase in time or errors with more perceptual units.

Method

Materials. The stimuli were 48 black shapes, each formed by filling in cells of a 4 x 5 square grid. Three levels of stimulus complexity were used, with one, two, or three perceptual units. A perceptual unit was defined as a set of one or more contiguously filled (i.e., black) horizontal or vertical cells of the grid (i.e., the Gestalt Law of Good Continuation was used to define the perceptual units) or a symmetrical group of three filled cells that formed a corner (i.e., the Gestalt Law of Good Form was used to define the perceptual units). Cells were filled randomly with the constraints that segments of the two and three-unit stimuli were connected to one another by shared sides or corners of adjacent segments. The one-unit stimuli had a mean of 3.0 filled cells; the two-unit stimuli had a mean of 4.6 filled cells; and the three-unit stimuli had a mean of 5.0 filled cells. Each of the 48 target stimuli appeared once.

Half of the stimuli in each set were paired with an identical stimulus, which corresponded to the "yes" trials; the other half were paired with a stimulus that differed from the first by the addition or deletion of one grid square from the shape, which corresponded to the "no" trials. In the "no" trials, both shapes had the same level of complexity; the alterations occurred on the first, second, or third unit of the shape, but due to constraints imposed by shapes in the 4 x 5 grid, a three-unit stimulus never had the alteration on its second (middle) unit.

Procedure. At the beginning of testing (before any actual experiment), we trained the subjects to press the response keys. In this training session, the word "yes" or "no" appeared in the center of the screen, and the subjects simply pressed the corresponding key as quickly as possible. If the subjects made an error, the computer beeped. Each word appeared 32 times; the trials were organized into two blocks, each of which contained a roughly equal number of both words. The words were presented in a random order, except that the same word could not appear more than three times in a row. All subjects were able to perform this experiment virtually perfectly by the second block of trials.

In each trial of the Shape Comparison task, the subjects were first asked to study the initial shape of a pair until they could remember it. They then pressed the space bar, and the screen went blank. After a 1 s delay, the probe shape appeared. The subjects were asked to respond "yes" if the probe shape was identical to the first member of the pair, or "no" if it was different. The one-unit stimuli were presented before the two-unit stimuli, which in turn preceded the three-unit stimuli. A typical trial sequence is illustrated in Figure 3.

Insert Figure 3 About Here

The first stimulus of a pair was always presented in the center of the screen, but the second was displaced a distance equal to one row up or down or one column to the left or right. This displacement was used to prevent the subjects from remembering the location of units on the screen itself or using an afterimage to make a response.

Results and discussion

In this and all other experiments, a score for response times and a score for error rates was obtained for each subject. In this experiment, these scores were computed by subtracting the time or errors for one-unit stimuli from those for three-unit stimuli. Two *t* tests were then performed, comparing R.V.'s response time and error rate scores to those from the control subjects. A deficit was inferred if either of R.V.'s scores fell outside the normal range and the manipulation (increasing the number of perceptual units) caused a monotonic increase in that dependent measure. Thus, although we computed scores using only the extreme values of the manipulation, the intermediate value plays a valuable role: If the manipulation was in fact progressively taxing a specific subsystem, then performance should be progressively impaired with greater values of the manipulation. (Note, however, that we do not know the underlying psychological scale affected by our manipulation, and hence cannot predict a linear increase, or any other quantitative relation, between the different values of the manipulation.)

As is illustrated in Figure 4, R.V. required progressively more time to respond to the more complex stimuli; in contrast, the normal control subjects did not show such an increase. Not surprisingly, R.V.'s response time score was dramatically different from those of the control subjects, $t(7) = 10.62$, $p < .001$. R.V. also made relatively more errors for the complex stimuli than did the control subjects, $t(7) = 9.96$, $p < .001$.

Recall that response times and errors trade off against each other: If we had urged R.V. to respond more quickly (e.g., by imposing a deadline), his error rates would have increased even more (e.g., see Luce, 1986). Thus, a deficit can be reflected by either score. The instructions emphasized responding as quickly and accurately as possible, and R.V. was particularly concerned about responding accurately (indeed, he often kept a running score of the number of errors he thought he had made!).

Insert Figure 4 About Here

In short, we found that R.V. did have a deficit in a simple, non-semantic visual task: He required progressively more time to respond to more complex stimuli, whereas normal subjects did not. Although on the surface the task seems very simple, from the point of view of Kosslyn et al.'s analysis it is remarkably complex. Indeed, their analysis of the subsystems of high-level vision leads us to consider 10 distinct possible causes of this deficit. Furthermore, these causes are not exclusive; any or all of them could be involved here, singly or in combination. Given the extensive region of damage and hypometabolism in R.V.'s brain, we cannot rule out a priori any of the following possible functional impairments.

1. *Visual buffer.* The visual buffer is a set of retinotopically mapped areas in prestriate cortex. This structure organizes edge fragments and regions into figure versus ground. This structure could have regions of hypometabolism or scotoma. If so, then the more complex the figure, the more likely it would be to fall on a dysfunctional portion of the buffer, making an eye movement necessary. Hence, more time would be required to evaluate the more complex stimuli.

2. *Attention window.* The attention window operates within the topographically mapped areas of the visual buffer, surrounding material in a specific region and sending this information further into the system for additional processing (cf. Treisman, 1990). If the attention window were restricted in scope, so that only part of the figure could be taken in at once, more complex figures would require one to move the attention window. Hence, more complex figures would require more time to examine than simple ones.

R.V.'s deficit also could reflect damage to subsystems in the ventral system, as follows.

3. *Preprocessing.* The preprocessing subsystem inferred by Kosslyn et al. (1990) extracts collinear edge fragments, symmetrical edges, points of intersecting edges, and other "nonaccidental" properties that are useful for recognizing objects when they appear at different sizes or orientations (for a good summary of the nonaccidental properties originally proposed by Lowe, 1987a, b, see Biederman, 1987). This subsystem may be impaired so that only a limited number of nonaccidental properties can be

extracted at a time. If so, then the presence of the grid lines may have overloaded the preprocessing subsystem so that it could not encode all of the edges and regions of the figure at one time. Hence, the grid lines may have forced this subsystem to encode one perceptual unit at a time, and so stimuli with more units would require more time to encode.

4. *Pattern activation.* The pattern activation subsystem is a modality-specific visual memory that stores representations of shapes. Input from the preprocessing subsystem selectively activated stored patterns. Two types of damage to this subsystem could produce the deficit: First, the pattern activation subsystem could be damaged so that it is difficult to store a visual representation of the shape of the first stimulus of the pair. If so, then the more complex the figure, the more degraded the stored representation would be, and hence the more time would be required to compare it to a probe stimulus. Second, the stored representations of shape may be intact, but this subsystem could be damaged so that the comparison process is impaired. In this case, the more complex the input from the preprocessing subsystem, the more time would be required to compare the probe stimulus to stored representations. Both functions could, of course, be impaired.

So far we have assumed that R.V.'s brain encoded the stimuli as shapes, using subsystems of the ventral system. However, it is possible that R.V. also encoded the shapes as sets of filled cell locations, using subsystems of the dorsal system. Indeed, when interviewed afterwards, R.V. claimed that he tried to remember the patterns by noting which individual cells were filled. If the ventral system were impaired, it may have encoded shapes slowly or poorly; this conjecture is consistent with the region of hypometabolism in the left occipital-temporal region. If so, because the dorsal system was relatively intact, its output could be used to make the judgment. Because the ventral and dorsal systems operate in parallel, the subject's performance will reflect properties of one or the other set of processes, depending on which system produces useful output first.

If R.V.'s decisions were based on such encodings, then his response times would be sensitive to variables that affect the ease of encoding locations, whereas the control subjects would produce the responses via the ventral system, which was not sensitive to these variables. Several factors could cause the deficit if the dorsal system were awry.

5. *Spatiotopic mapping.* If R.V.'s responses reflect processing in the dorsal system, he may have had impaired performance because the second stimulus of a pair was displaced. The spatiotopic mapping subsystem locates objects relative to the body or another object, not the retina. If this subsystem were impaired, it would require a relatively long time to register the location of each segment, and hence the more complex the stimulus, the more time would be needed to encode it. We did not expect this hypothesis to be borne out, given that the parietal lobes are intact; nevertheless, we felt it important not to succumb to a "confirmation bias," and explicitly checked implausible hypotheses. Two other deficits in the dorsal system were plausible, however, as noted below:

6. *Categorical spatial relations encoding.* If the pattern were encoded as a configuration of locations, they may have been specified relative to the grid itself. The categorical spatial relations encoding subsystem assigns relative positions to categories, such as "top," "leftmost," or "connected to." These representations are efficient for encoding locations of filled cells in a grid. The categorical spatial relations subsystem itself is posited to be in the posterior parietal lobe (on the left side, as is evident in left/right confusions following damage to the left angular gyrus; see De Renzi, 1982); hence, we do not expect this subsystem to be impaired. However, the output from this subsystem projects to the frontal lobes; if these connections are damaged, more time could be required to encode more complex stimuli. Thus, the impaired performance thus may reflect damage to the connections from the categorical spatial relations encoding subsystem as well as damage to the ventral system.

7. *Coordinate spatial relations encoding.* The locations of the filled cells also could be specified using metric distances, and the coordinate spatial relations encoding subsystem encodes metric distances (for a discussion of the distinction between categorical and coordinate spatial relations representations, see Kosslyn, Koenig, Cave, Barrett, Tang, & Gabrieli, 1989). The same/different decision could be based on the output from the coordinate spatial relations subsystem if the output from the categorical spatial relations encoding subsystem were sufficiently degraded. This seemed plausible, given that the lesion disrupted processing in the left hemisphere, and the coordinate spatial relations subsystem is more effective in the right hemisphere (Kosslyn et al., 1989). If so, then the deficit would not be due to this subsystem's being disrupted. However, our hypothesized anatomical localization could be awry;

thus, we thought it was important to discover whether R.V. could encode metric spatial information properly. If not, then his deficit could arise if both spatial relations subsystems were impaired and the decisions were based on encodings of shapes as sets of locations.

8. *Associative memory.* The output from the dorsal and ventral systems must be sent to an associative memory. The mere fact that we can recall the locations of objects on our desk indicates that object properties and spatial properties are associated in memory. In this task, the associative memory may store the set of locations of filled cells of the first stimulus, and compare these locations to those of the second stimulus. Goldman-Rakic (1987) describes a spatial short-term memory in area 46 of the frontal lobes, which appears to be serve this function. If this subsystem is damaged, R.V. may have trouble storing the locations of the filled cells, particularly on the contralateral side. The more complex the stimuli, the more they would tax the impaired memory system, resulting in increased response times. Furthermore, the decision produced on any given trial must be mediated by information in associative memory, which is critical for understanding the task and for evaluating the products of prior processing appropriately.

9. *Top-down processing.* The visual system does not passively wait for new information; rather, hypotheses are formed and actively tested (e.g., see Gregory, 1970). Such top-down processing is particularly likely if a subtle discrimination is necessary, and one must take a "second look" to obtain enough information to perform the task. Because the "different" probe stimuli (on "no" trials) were relatively similar to the initial study stimuli in the Shape Comparison task, such "second looks" may have been used at least some of the time. It is possible that top-down processing was used more often with more complex stimuli because they are more difficult to represent fully in a single encoding. If so, then an increase in time with complexity may reflect damage to the categorical or coordinate property lookup subsystems or to the categorical-coordinate conversion subsystem.

10. *Attention shifting.* A set of subsystems is necessary to shift attention over a stimulus. Posner, Inhoff, Friedrich, and Cohen (1987) hypothesize that subcortical structures, specifically the superior colliculus and thalamus, are used to shift attention and engage it, respectively. It is possible that critical connections from these structures were impaired. Thus, although even normal people may examine the stimuli a part at a time, they may be able to shift their attention (i.e., scan over it) much faster than R.V. If R.V. has an impaired ability to shift attention, he may require more time to examine more complex stimuli.

In addition, it is possible that R.V. simply tired by the time the three-unit stimuli were presented. This post-hoc explanation is not very convincing, given the relatively few trials (indeed, one could have just as easily predicted decreased times with practice). Nevertheless, we will address this possibility in the course of ruling out various other interpretations.

Experiment II: Short-Term Memory Control

We begin by asking broadly whether the deficit reflects impaired memory for the first stimulus, encoded either as a shape or as a set of locations. In this experiment, the subjects saw one of the stimuli used in the Shape Comparison experiment along with an X mark, and simply indicated whether the X fell on or off the shape. If R.V.'s deficit occurred solely because he has trouble remembering the first stimulus of a pair, then it should not be evident in this experiment. As before, the *manipulation* was the number of perceptual units; by varying the number of perceptual units, we taxed the subsystems that encode the pattern. In contrast to the Shape Comparison task, this task does not require remembering a pattern; hence, an impaired pattern activation subsystem or associative memory should not cause a deficit in this task. The *score* was the amount of increase in time or errors with more perceptual units.

Method

Materials. The first stimulus of the pairs used in the Shape Comparison experiment were used here. In this experiment, however, the patterns were presented in a light gray tone instead of the solid black used before; the gray tone was necessary so that the black X probes would be visible on a "yes" trial (when they appeared on the pattern). The grid lines inside the segments of the target were removed so that the target still appeared as a solid shape within the grid. As in the Shape Comparison experiment, there were 48 trials, with each shape appearing just once. Half of the trials were "yes" trials, in which the X fell on the shape, whereas the other half were "no" trials, in which

the X fell in a cell adjacent to the shape. As before, the one-unit stimuli were presented before the two-unit stimuli, which in turn were presented before the three-unit stimuli.

Procedure. The beginning of each new trial was announced by an exclamation mark on the screen. The subjects pressed the space bar, and 1 s later the stimulus (pattern-with-X-mark) appeared. This stimulus remained visible until the subjects pressed one of the two response keys, at which point the exclamation mark appeared again to signal the beginning of a new trial.

Results and discussion

The scores were computed as in the Shape Comparison experiment, and the results are illustrated in Figure 5. As is evident, the time required for R.V. to respond again increased for increasingly complex stimuli and this increase was not present in the data from the control subjects, $t(7) = 15.01$, $p < .001$. In this case, there was no difference between the error rate scores for R.V. and the control subjects, $t < 1$.

Insert Figure 5 About Here

Clearly, the deficit observed in the Shape Comparison experiment was not due solely to impaired short-term memory. Even when we eliminated the memory component, a deficit was still evident. Furthermore, the fact that a deficit persisted even when the stimulus was not moved on the screen suggests that the deficit in the Shape Comparison experiment was not due solely to an impaired spatiotopic mapping subsystem. An impaired spatiotopic mapping subsystem would affect processing only when retinotopic representations could not be used to perform the task; the present task could in fact have been performed with such representations.

Experiment III: Pattern Activation Encoding

The previous results suggest that R.V.'s problem with the Shape Comparison experiment was not solely a consequence of impaired short-term memory. However, the increase in time with complexity in the Shape Comparison experiment was about twice that in Experiment II, which might suggest that *impaired memory* contributed to the deficit in the Shape Comparison experiment; R.V. might have trouble encoding new shapes into the pattern activation subsystem. This possibility was evaluated in this experiment. We asked the subjects to study a shape in a set of brackets, with the internal grid lines removed. Thus, they could not encode the shape as a set of filled locations, using the dorsal system, and were forced to encode it as a shape. After studying the shape, it was removed, and the subjects were forced to remember the shape. An X mark was then presented as a probe within a set of brackets, and the subjects decided whether this probe occupied a spot that previously was covered by the shape. The *manipulation* and *score* used here were the same as in the previous two experiments.

Method

Materials. The shapes and X probes used in the Short-Term Memory Control experiment were used here. In this case, however, two stimuli were presented on each trial; one containing only a shape, and the other only an X mark. In both cases, the internal grid lines were removed and only the four corners of the external frame were retained, as is illustrated in Figure 6. In addition, different X probes were paired with the shapes, and the stimuli were presented in a different order than in the previous experiment; however, as before, all of the one-unit stimuli were presented before all of the two-unit stimuli, which in turn were presented before all of the three-unit stimuli, and half the trials at each level of complexity included X's that could be superimposed on the shape ("yes" trials), and half included X's that fell adjacent to the shape ("no" trials).

Insert Figure 6 About Here

The shapes were presented in the center of the screen, but the probes were displaced one row's width up or down or one column's width length left or right; the displacement was used to prevent the subjects from remembering locations on the screen itself or using an afterimage to make a response.

Procedure. To announce the beginning of a trial, an exclamation point appeared in the center of the screen and disappeared when the subject pressed the space bar. A black shape in a set of brackets then appeared; when the subjects had memorized the shape, they pressed the space bar. The screen then went blank and remained so for 2500 ms, at which point an X probe inside an empty set of brackets

appeared. The subjects were to decide whether the X would have fallen on the shape were it still present and the brackets were aligned. If so, they were to respond "yes"; if not, they were to respond "no". In all other respects, the procedure was like that used in the Short-Term Memory Control experiment.

Results and discussion

The results were strikingly different from the previously described two experiments: As is illustrated in Figure 7, there was no difference between the control subjects and R.V. in the response time scores, $t(7) = 1.28$, $p > .1$, and R.V. actually did better than the control subjects in the error rate scores, $t(7) = -5.85$, $p < .01$.

Insert Figure 7 About Here

Thus, we have good evidence that R.V. can effectively store shapes in the pattern activation (i.e., modality-specific visual memory) subsystem. When we used stimuli that could not easily be encoded as sets of locations, we found that he could indeed remember and compare complex shapes as well as simple ones. These results indicate that the capacity limitations of the pattern activation subsystem did not contribute to the results of the Shape Comparison experiment. Furthermore, they allow us to eliminate the hypothesis that R.V.'s increased times with more complex stimuli merely reflect increased fatigue. This experiment had the same number of trials as the previous one, and the stimuli were presented in order of increasing complexity, yet there was no evidence of increased time with increasing complexity. Indeed, R.V. was very vigorous throughout testing and showed no signs of flagging interest or ability.

Experiment IV: Preprocessing Overload

It is possible that R.V.'s deficits in the Shape Comparison and Short-Term Memory Control experiments were caused by an impaired preprocessing subsystem. The preprocessing subsystem is posited to extract the aspects of shape that are invariant over a wide range of different projections of the object (Lowe, 1987a, b). If this subsystem has been damaged, it may fail to encode enough of these characteristics to recognize the shape immediately. A damaged preprocessing subsystem may become overloaded by the lines of the grid, which would interfere with the encoding of the nonaccidental properties of the shape itself. This effect would be more severe for more complex shapes because they have additional nonaccidental properties, and so present an even greater load to an already-taxed preprocessing subsystem than simple shapes. However, despite the interference from the grid lines, the ventral system may still operate more efficiently than the dorsal system, and so the response would reflect this impairment.

In this experiment, the subjects again merely indicated whether an X mark was on a figure. Now, however, the figure was presented in an empty frame. If the grid lines were overloading the preprocessing subsystem, then R.V. should not show increases in time with complexity in this experiment. As before, the *manipulation* was the number of perceptual units in the figure (1, 2, or 3) and the *score* was the increase in time or errors with perceptual units.

Method

Materials. The materials from the Short-Term Memory Control experiment were used here, except that the internal grid lines were removed. The stimuli were left with only the outside four corners (brackets) of the original grid.

Procedure. The procedure was identical to the Short-Term Memory Control experiment in which grids were used.

Results and discussion

These results were analyzed as in the previous experiment, and are illustrated in Figure 8. As is evident, the increase in time with complexity was eliminated when the grid lines were removed, $t(7) = 1.85$, $p > .25$, and there was no difference between R.V.'s error rate score and those of the control subjects, $t < 1$.

Insert Figure 8 About Here

Thus, we have evidence for one source of R.V.'s impaired performance in the Shape Comparison experiment. When the grid lines were present in this task, we found an increase in time with complexity; when they were removed, there was no such increase. This finding is consistent with the idea that the preprocessing subsystem was overloaded when the grid lines were present. This inference is also consistent with the fact that PET scanning indicated hypometabolism in the occipital-temporal area, which is where the preprocessing subsystem is hypothesized to be localized (Kosslyn et al., 1990).

Experiment V: Ventral Shape Comparison

The results described so far suggest that the grid lines played a critical role in the observed deficits in the Shape Comparison and Short-Term Memory Control experiments. If so, then the deficit in the Shape Comparison experiment should be eliminated simply by eliminating the grid lines. In this case, the preprocessing subsystem should be less taxed. In all other respects, this experiment was identical to the Shape Comparison experiment.

Method

Materials. The stimuli for the Shape Comparison experiment were used here, except that the grid lines were removed from all the stimuli, leaving only the outside four corners (brackets) of the original grid and the black shape.

Procedure. The procedure was identical to that of the Shape Comparison experiment.

Results and discussion

The data were analyzed as in the Shape Comparison experiment, and the results are presented in Figure 9. As is evident, removing the grid lines had the expected effect: We no longer found increased times with increasing complexity, and these results were no different from those of the control subjects, $t(7) = -1.02, p > .25$; similarly, there was no difference between R.V.'s error rate score and those from the control subjects, $t(7) = 1.87, p > .1$.

Insert Figure 9 About Here

These findings, then, buttress our inference that the grid lines were at the root of the observed deficit. However, we noted earlier that by eliminating the grid lines, we also made it difficult--if not impossible--to encode the patterns as sets of filled locations. Thus, it is possible that the reduced metabolism in the ventral system, evident in R.V.'s PET scan, impaired the preprocessing subsystem. As a consequence, the dorsal system may often have produced a representation of the locations of filled cells more quickly than the ventral system produced a representation of the shape. In this case, the output from the dorsal system would actually underlie his response. If so, then the impaired performance in the Shape Comparison and Short-Term Memory Control experiments would reflect limitations of the location-encoding system, which ultimately dictated the pattern of response times. Thus, the results described so far do not eliminate possible difficulties in the dorsal system.

Experiment VI. Categorical Spatial Relations Encoding

The fact that R.V.'s performance was impaired even when the display was not displaced (in the Short-Term Memory Control experiment) suggests that a damaged spatiotopic mapping subsystem was not the root of his problem. However, it is possible that R.V. had trouble representing the locations of filled cells in the grid. As noted earlier, when interviewed afterwards, R.V. claimed that he tried to remember the stimuli in grids by noting the location of each filled cell. Given that the grid provides a convenient framework for using categorical spatial relations representations, we considered the possibility that R.V. had trouble encoding patterns in grids because he did not encode categorical spatial relations effectively. To explore this hypothesis, we showed the subjects a horizontal bar and an X, and asked them to decide whether the X was above or below the bar. The location of the bar and the X moved from trial to trial, so that the subjects had to encode a spatial relation; they could not simply look at a part of the screen to make the decision. The *manipulation* was the distance of the X from the bar; it was either very close to the bar or over 2 cm from it. The *score* was the increase in time and errors when the X was close to the bar compared to when it was farther from the bar.

Method

Materials. In this experiment, all stimuli contained a bar and an X. The bar was a horizontal segment, of the same size as four contiguous cells in the grids used in previously described experiments.

This bar was placed roughly in the center of an elongated set of brackets, as is illustrated in Figure 10. The bar could be located in one of two positions, one of which was the bar's height above the other. The experiment included 64 trials; for each bar, the X probe was positioned in one of 32 relative locations. The 32 probes in each set were evenly divided so that 16 were above the bar, and 16 were below it. For each of these categories, 8 X's were within .5 inch of the bar, presenting a difficult discrimination task, and 8 X's were outside .5 inch of the bar, presenting an easy discrimination task.

As will be discussed shortly, the same stimuli were also used in a metric distance judgment task. Thus, we also counterbalanced the difficulty of that decision with the other variables. Of the 16 X's per bar that were inside the invisible .5 inch boundary (8 above the bar, and 8 below the bar), 8 of them (4 above, 4 below) were close to the .5 inch boundary (which will correspond to "difficult" metric discriminations) and 8 (4 above, 4 below) were relatively far from the boundary ("easy" metric discriminations); of the 16 X's per bar that were outside the invisible .5 inch boundary (8 above the bar, 8 below the bar), half were close to the boundary ("difficult" discriminations) and half were relatively far from it ("easy" discriminations).

The X's were placed in four locations horizontally relative to the bar (equivalent to being in the four columns of the grid used in the other experiments). A Latin Square design was used so that every stimulus variation occurred equally often in each of the horizontal positions.

Insert Figure 10 About Here

The 64 trials were divided into two blocks of 32; each block was counterbalanced with a Latin square design for the variables above/below, easy/difficult, left/right half of bracket, and bar location. The trials were randomized with the constraint that none of the conditions of the following variables were repeated more than three time in a row: bar location, above/below position, easy/difficult discrimination, left/right location of X relative to the bar, and central/peripheral location of X along bar.

Procedure. As usual, each trial began with an exclamation point, which remained on the screen until the subjects pressed the space bar. After a 1 s delay, during which the screen was blank, the stimulus appeared. The subjects decided whether the X was above or below the bar; if above, they pressed the "yes/above" key; and if below, they pressed the "no/below" key. The response keys were labeled in this way to remind the subjects of their function. Immediately after the subjects responded, the exclamation point returned and a new trial began.

Results and discussion

R.V. had a larger response time score than the control subjects, $t(7) = 4.62$, $p < .01$, but there was no difference in the error scores, $t < 1$. Thus, there was evidence of a deficit in R.V.'s ability to encode at least one categorical spatial relations representation, above/below.

One might question whether we had had reason to expect any deficit in spatial relations encoding at all, given that R.V.'s parietal lobes were spared by the damage. Recall, however, that Goldman-Rakic (1987) has found that the frontal lobes are critically involved in processing spatial information, and that R.V. has a frontal lobe lesion in an area that may be the homolog of that studied by Goldman-Rakic. The left-hemisphere advantage in processing categorical spatial relations (Hellige & Michimata, 1990; Kosslyn et al., 1989) is consistent with this deficit, given that R.V. had a left-hemisphere frontal lobe lesion. Damage to the superior longitudinal fasciculus might suggest that the frontal lobe was not able to use information from the left parietal lobe as effectively as it could prior to the stroke.

Experiment VII. Coordinate Spatial Relations Encoding

This experiment utilized the same materials used in the categorical spatial relations encoding experiment, except that now the subjects were asked whether the X fell within .5 inches of the bar (and ignored whether the X was above or below the bar). The *manipulation* was the difficulty of the discrimination; when the X was between .4 and .6 inches, the discrimination was difficult, whereas when it was between .1 and .3 or between .7 and 1 inches, the discrimination was easy. The *score* was the increase in time and errors when the X was close to the criterion compared to when it was farther. If R.V. has a deficit in encoding metric information, it should be exacerbated in the more difficult condition (i.e., when the X was close to the criterion).

Method

Materials. The materials used in this experiment were identical to those used in the Categorical Spatial Relations Encoding experiment.

Procedure. The procedure was similar to that of the Categorical Spatial Relations Encoding experiment. At the beginning of the task, however, two samples of a bar embedded in elongated brackets appeared on the screen. One sample had a horizontal dotted line drawn .5 inch from the top edge of the bar, and the other sample had a horizontal dotted line drawn .5 inch from the bottom edge of the bar. The subjects were asked to memorize how the half-inch distance looked in both samples on the screen. After 12 practice trials in the task, with feedback, the two samples returned to the screen, and the subjects were instructed to press the space bar when ready to begin the actual experiment.

The same stimulus sequence used in the previous experiment was used here, except that now the subjects were asked to decide whether each X was within a half-inch distance of the bar. If it was, they were to press the "yes/in" key; if it was not within the a half-inch distance, they were to press the "no/out" key. Again, the response keys were labeled in this way to remind the subjects of their functions.

Results and discussion

R.V. had a deficit in encoding coordinate spatial relations, as indicated by a difference in the response time scores, $t(7) = 7.77$, $p < .001$; there was no difference in the error scores, however, $t(7) = -1.26$, $p > .1$.

This finding was somewhat surprising, given the evidence that categorical spatial relations are encoded more efficiently in the left cerebral hemisphere and coordinate spatial relations representations are encoded more efficiently in the right cerebral hemisphere. We will return to this result after we have considered the findings from all of our experiments.

Experiment VIII: Location Associative Memory

R.V. has a lesion near what may be the human homolog to area 46. Thus, it seemed possible that he may have a specific deficit in short-term memory for location. To explore this possibility, we asked the subjects to study either 2 or 4 gray blocks within a set of brackets. The subjects memorized the location of the blocks, and then the blocks were removed for 1 s, at which point an X appeared. The subjects were to decide whether the X was in a location previously occupied by a block. The *manipulation* here was the number of blocks, and we expected increased times and errors with additional blocks if any of the subsystems involved in representing location were awry. The *score* was the increase in time or errors for 2 versus 4 blocks.

Method

Materials. The stimuli consisted of a set of four brackets, placed at the corners of an invisible 4 x 5 grid. Within the brackets were either two or four gray blocks, each of which was the size of a cell in the 4 x 5 grid. The blocks were separated by at least one block's width from each other; this prevented subjects from merging two or more of the blocks into a single perceptual unit, and thus forced them to remember the location of each block as separate unit. The blocks appeared equally often in the four quadrants in both the two-block and four-block displays.

The probe stimuli were a set of brackets containing a single X mark. On "yes" trials, the X fell in a position that was previously occupied by a block; on "no" trials, the X fell adjacent to a position that was previously occupied by a block. The X probes appeared equally often in the left and right halves of the brackets, and the "no" probes fell equally often above, below, to the left, and to the right of locations that contained blocks. Each stimulus was presented twice, although the same stimulus was

never presented on consecutive trials. On one presentation, the stimulus was followed by a "yes" probe, and on the other it was followed by a "no" probe. The entire experiment consisted of 48 trials.

Procedure. At the beginning of each trial, an exclamation point appeared. The subjects pressed the space bar, and the exclamation point was replaced by a set of brackets containing either two or four gray blocks. After studying the blocks, the subjects pressed the space bar and the stimulus disappeared; 1 s later a set of brackets with an X probe appeared. The subjects were asked to respond "yes" if the X fell in a location that previously had held a block, and "no" if it fell in a location that previously had been empty.

Results and discussion

The results are illustrated in Figure 11. As is evident, R.V. did in fact have a deficit in this experiment in the response time score, $t(7) = 11.01$, $p < .001$. There was, however, a trend for the controls to have larger error scores than R.V., $t(7) = -2.17$, $p > .05$. We also analyzed R.V.'s relative performance for probes in the left versus right halves of the display, and found no differences, $t < 1$ for both response times and errors.

Insert Figure 11 About Here

Thus, we have evidence that R.V. did have a deficit in his ability to store information about location. This is remarkable given that he only had to remember the locations for 1 s (the same time as in the original Shape Comparison experiment). However, in additional analyses we did not find the human analog to Goldman-Rakic's finding that the deficit was for locations in the field contralateral to the lesion; R.V. did not have particular trouble retaining information about location on the right side of space. The stimuli only subtended 3.1° of visual angle horizontally, however, and this may not have been enough to tax the contralateral spatial memory.

Experiment IX: Preprocessing Followup

We have evidence, then, that R.V. has a deficit both in his ability to extract nonaccidental properties (i.e., in his preprocessing subsystem) and in his ability to encode and retain metric spatial information. We have assumed that the dorsal system would be used in the Shape Comparison task only if the ventral system were impaired, so that it "lost" the race to send output downstream. This experiment was designed to provide converging evidence for such a deficit in the ventral system. It was identical to the preprocessing overload experiment, except that random line fragments were placed over the stimuli. These fragments were irregularly positioned, and sometimes intersected with one another or with the gray stimulus pattern. The fragments did not form distinct cells, eliminating the option to encode the pattern as a set of locations in a grid. Thus, these stimuli forced the subjects to encode the patterns as shapes, and should have made the task relatively difficult if the preprocessing subsystem were impaired (by taxing or overloading the subsystem with lines and intersections that are irrelevant to the task). The *manipulation and score* were the same as in the original Shape Comparison experiment, namely the number of perceptual units and the effect of increased units on performance.

Method

Materials. The stimuli were constructed by adding four vertical and three horizontal thin lines of varying lengths to each stimulus from the Pattern Activation Encoding experiment (Experiment III). Although the same seven fragments were added to every stimulus, the lines were positioned differently for each stimulus, resulting in a different configuration of overlapping segments for each. The stimuli were constructed in this way in order to prevent subjects from adjusting to a particular pattern of line fragments, while keeping the total number of added segments and the length of the segments constant across all stimuli. The stimuli were presented in the same order as in the Pattern Activation Encoding experiment.

Procedure. The procedure was identical to that used in the Pattern Activation Encoding experiment.

Results and discussion

R.V.'s times did in fact increase with increasing complexity relative to those of the control subjects, $t(7) = 1.51$, $p < .05$ (using a one-tailed test, which is justified given that we predicted the direction of the difference). Note also in Figure 12 that there was a monotonic increase in times with

complexity, as expected if this variable were increasingly taxing the subsystem. In addition, R.V.'s error score was larger than the control subjects', $t(7) = 8.68$, $p < .01$.

As is evident in Figure 12, the effect was not as dramatic as before, which may be a consequence of at least three factors: First, fewer line segments appeared here than appeared in the grid. Hence, this display may not have taxed the preprocessing subsystem as much as the grid. Second, because the lines did not define discrete locations, the location-based strategy could not be used. It is possible that although this processing resulted in faster overall times, and hence "won" the race, it still displayed an abnormal sensitivity to increasing complexity. Third, and most mundane, we must note that this experiment was administered 6 months after the initial ones. Thus, R.V. could simply have improved in the meantime. This seems unlikely, however, because his mild reading problems, which may be a result of the preprocessing limitations observed here, had not improved.

In any event, the most important finding here is that R.V. did exhibit impaired processing when the preprocessing subsystem was taxed by spurious lines, even when these lines did not encourage location-based encoding. Thus, we have evidence that the ventral system was indeed impaired.

Insert Figure 12 About Here

Experiment X: Location Top-Down Search

Although we have evidence that R.V. has impaired preprocessing, spatial relations encoding, and location associative memory subsystems, we have not exhausted the possibilities. It was possible that at least some of his problem is in taking "second looks" at patterns when comparing them. Thus, we conducted a series of experiments to examine how well R.V. could use stored information to direct his attention.

We designed two sets of experiments to examine the possibility that R.V. has a deficit in using stored information to guide top-down search. In one, we examined his ability to use stored information to direct attention to a particular location, using a task that was previously employed by Kosslyn, Cave, Provost and Von Gierke (1988) to study visual mental imagery. We were not interested in its imagery components, but rather in the requirement that one access memory to determine where a part of a letter should be placed.

In this task, the subjects first studied upper case block letters which appear within four brackets. Each block letter was associated with a lower case, cursive cue. In the task, the cue was presented briefly in the center of the screen, and then was replaced by a set of brackets containing only a single X mark. The subjects were asked to decide whether the X would have fallen on the corresponding block letter if it were present within the brackets as studied previously. Kosslyn et al. (1988) found that subjects required more time for letters that had more segments, which suggests that at least some of the subsystems used to generate images must work harder to image letters that have more segments. Our *manipulation* was the number of segments in the block letter, and the *score* was the increase in time and errors with more complex block letters.

Method

Materials. Block versions of four, three-segment "simple" letters (C, F, H, U) and four, five-segment "complex" letters (J, P, S, G) were formed by filling in cells of 4 x 5 grids. Each letter was then presented within four brackets, centered on the screen; the brackets corresponded to the corners of the grids used previously, with all other lines removed.

A cursive lower case version of each letter was paired with the corresponding block letter; the cues were presented in the center of the screen. Each cue was paired with four brackets stimuli, each containing a single X mark. For two of the trials, the X would have fallen on the corresponding block letter were it within the brackets; for the other two, it would have fallen adjacent to a segment of the block letter. Thus, there were a total of 32 trials in this experiment. Two additional letters, L and O, were used in practice trials only.

We also considered the order in which the segments of the block letters were typically drawn (see Kosslyn et al., 1988, for details). For each letter, one "yes" and one "no" X probe was placed on or near to a segment that was drawn early in the sequence, and one "yes" and one "no" X probe was placed on or near a segment that was drawn late in the sequence.

Procedure. The subjects first participated in a task to teach them the cue-block letter association. They began by reviewing the block letters and the corresponding cursive letters, pressing the space bar to see the next paired cursive and block letter. The subjects studied each pair as long as they wished. After seeing three randomized sets of the letters, the subjects were then presented with the cursive letters one at a time. They were given a black, thick-tipped marker and asked to draw the corresponding block letter on paper; the paper contained empty sets of brackets that were the same size as those on the screen. The subjects were reminded to place the block letter correctly inside the brackets. After the subjects had drawn all 10 letters, the experimenter checked the drawings for accuracy. If the subjects drew any of the letters incorrectly, the letter pairs were displayed again one at a time on the screen. When the subjects had correctly drawn all the letters of the set, placing them properly inside the brackets, the experimenter stopped the letter learning session.

Another experiment, not described here (a perceptual control that turned out to be unnecessary, and hence an unnecessary burden on the reader to describe), was then conducted. Following this, the present experiment was conducted. It began with an exclamation point, which remained in the center of the screen until the subjects pressed the space bar; at this point the screen became blank, and 500 ms later, a lower case cursive letter appeared in the center of the screen for 500 ms. This letter was a cue to image the corresponding block letter as it would appear in the set of brackets. A blank screen was then presented for 500 ms before an X appeared in an empty set of brackets. The subjects were asked to respond "yes" if the X occupied a spot in the brackets that would be occupied by the block version of the cued letter, and "no" if the X occupied a spot in the brackets that would not be occupied by the block letter. After the response, the exclamation point returned to the screen to signal the beginning of the next trial.

Results and discussion

We found no evidence for a deficit in this task, $t(7) = -1.27$, $p > 1$ for the response time score, and $t < 1$ for the error rate score. Thus, R.V. has no difficulty accessing the stored locations of individual segments of a shape. The intact performance here suggests that the processes that access stored information about the spatial structure of objects are not specific to the left frontal lobe. It is possible that the right hemisphere stores such information, which can be used in this task (cf. Kosslyn, in press).

Experiment XI: Shape Top-Down Search

We inferred that R.V. had no trouble accessing the specifics of the structural description of a shape. This experiment was designed to assess the ease of accessing stored information about the shapes of parts of objects. The subjects were shown a bar, cued with a cursive letter, and then shown an incomplete block letter. The question was, when the bar is added to the incomplete letter, do they form the block letter corresponding to the cursive cue? To perform the task, the subjects must access the stored representation of the proper block letter, compare the stimulus to that block letter and note the missing segment. They then must determine whether the previously-displayed bar would complete the block letter. This is a complex task. However, the *manipulation* again was the number of segments of the stored representation. If there was a deficit in accessing the stored description of how segments are arranged to form the block letter, then this manipulation should affect the difficulty of the experiment. The score was the same as in the previous experiment.

Method

Materials. Each trial included three stimuli. The first consisted of a horizontal or vertical black bar to study; these bars were 2, 3, or 4 "grid cells" long. For the experiment as a whole, there were equal numbers of horizontal and vertical black bars studied, and these bars were equally distributed across the "yes" and "no" trials and were almost equally distributed for simple and complex letters (the simple "no" trials had one too many vertical black bars whereas the complex "no" trials had one too many horizontal black bars).

The second set of stimuli were lower case cursive cues, which appeared above the bar. The same eight letter cues used in the previously described task were used again here. The remaining two letters were used as practice stimuli, also as before.

The third set of stimuli consisted of block letters with one segment missing. On "yes" trials, the subjects were given a black bar that completed the block letter that was cued. On half of the "no" trials, the bars were the wrong length (in spite of correct orientation) to complete the cued, incomplete block letter; on the other half of the "no" trials, the bars completed the block fragment to form an incorrect

block letter (i.e., the bar completed a block letter that was not cued). This kind of "no" trial forced the subjects to pay attention to the cue, and it ensured that they accessed information about shape stored in associative memory. Each of the 8 letters appeared once before any of the letters was repeated. In addition, all the cursive cues appeared once before any single cue was repeated. The experiment was presented in two blocks. Both blocks were balanced for the number of trials for the orientation of studied black bar, letter complexity, near/far location of the missing segment, and response.

Procedure. This experiment was conducted after the previously described one, during the same testing session, and so no special training was necessary to familiarize the subjects with the appearance of the block letters or their cursive cues. A test trial began with an exclamation point, which appeared for 500 ms in the center of the screen, after which a horizontal or vertical black bar appeared in the lower part of the screen. When the subjects felt that they had memorized the size and orientation of the bar, they pressed the space bar. The horizontal or vertical bar remained on the screen, and a centered asterisk appeared in the center of the screen. After 500 ms, the asterisk was replaced by the cursive cue. (The black bar was still present beneath it). The cue stayed on the screen for only 500 ms, at which point both the cue and the black bar disappeared and were replaced immediately by an incomplete block letter inside brackets. The subjects were asked to decide whether the black bar they had just studied would complete the block letter that was paired with the cursive cue. If so, they were to press the "yes" key; if not, they were to press the "no" key. A typical trial sequence is illustrated in Figure 13. After the subjects responded, a new trial began.

Insert Figure 13 About Here

Results and discussion

R.V. did not have a deficit in this task, $t < 1$ for both the response time and error rate scores. This result is consistent with the findings from the previous experiment. Once he has identified the shape, R.V. can access information about the arrangement of the individual segments.

Experiment XII. Scanning

When taking "second looks" one uses stored information to help scan over an object. It was possible that grid lines impaired R.V.'s scanning, and thus he required more time than the control subjects when more segments had to be searched. Thus, we assessed his scanning ability. A donut-shaped grid was presented with 3 contiguous filled cells, and an arrow appeared within the central hole. The subjects were asked whether the arrow points at a filled cell. We expected the subjects to require more time to respond when they had to scan greater distances, and examined whether this increase was larger for R.V. than for the control subjects. Thus, the *manipulation* was the distance between the arrow and the grid (3 distances were used), and the *score* was the increase in time or errors with increasing distances.

Insert Figure 14 About Here

Method

Materials. As is illustrated in Figure 14, the stimuli consisted of a square ring of 20 cells, with 4 cells on each side and one at each corner. For each stimulus, three adjacent cells of the ring were filled in (blackened) at random and an arrow was positioned inside the square hole of the ring. On "yes" trials, the arrow pointed to the center of a filled cell; on "no" trials, the arrow pointed to the center of a cell that was not filled but was adjacent to a filled cell. The arrow pointed in one of eight directions (North, Northeast, East, Southeast, South, Southwest, etc.) and could be near ($.08^\circ$ of visual angle), moderately far ($.76^\circ$), or far (1.4°) from the nearest edge of the square to which it was pointing. The arrow appeared equally often in the left and right halves of the ring, and pointed in each direction equally often. The experiment was divided into two parts, each with 48 trials. Both parts were counterbalanced for distance, direction, location of the arrow, and response using a Latin square design.

Procedure. At the beginning of each trial, an exclamation point appeared and remained until the subjects pressed the space bar. The exclamation point then disappeared, and 500 ms later a stimulus appeared. The subjects were told simply to indicate whether the arrow pointed to the center of a filled

cell; if so, they were to respond "yes"; if not, they were to respond "no." After the subjects responded, the exclamation point reappeared, signaling the start of the next trial.

Results and discussion

The results are presented in Figure 15. R.V. did in fact require more time to scan than did the control subjects, $t(7) = 4.30$, $p < .01$, but made relatively fewer errors for the longer distances than the control subjects, $t(7) = -7.40$, $p < .001$. This unfortunate speed-accuracy tradeoff makes it difficult to interpret these results.

Insert Figure 15 About Here

Experiment XIII: Scope of Attention Window

It was possible that R.V. had difficulty attending to larger regions of space. If so, he may have had a tendency to look at the complex shapes a segment at a time. This experiment was designed to discover whether the attention window had an abnormally restricted scope. The subjects studied four gray blocks that were positioned along the circumference of an invisible circle. On half the trials, an "X" mark appeared in two blocks on opposite sides of the circle, and on the other half of the trials only one X mark appeared. The subjects responded "yes" if both Xs were present, and "no" if only one was present. The *manipulation* was the diameter of the circle, which was one of two sizes. If the attention window size were restricted, such that it could not easily be enlarged to cover the entire area of the larger circle, then we should find impaired performance on trials with stimuli placed on the circumference of the larger circle. Hence, the *score* was the difference in times and errors for the two sizes.

Method

Materials. The stimuli for this experiment consisted of a set of four brackets that contained four filled gray squares (blocks). The blocks were arranged at 90° intervals along the circumference of an invisible circle. (The circle did not appear on the screen, and was merely used to help position the squares during construction of the stimuli.) The blocks were 75% of the size of the standard cell size used in the other experiments, which allowed us to have a larger difference in the distance among them. There were two types of stimuli: One type included blocks that were arranged along the circumference of a small circle (subtending 1.5° of visual angle), and the other had blocks that were arranged along the circumference of a large circle (subtending 3.0° of visual angle). Furthermore, although every stimulus contained four blocks at ninety-degree intervals along the circumference, the absolute positions of the blocks along the arc of the circle was varied. For example, a stimulus could have blocks at the 0°, 90°, 180°, and 270° positions along the circumference, or at the 36°, 126°, 216°, and 306° positions. There were five different positions of blocks along the arc of the circles, starting at 0°, 18°, 36°, 54°, and 72°. These five positions, together with the two sizes of the circle (large and small), allowed ten unique stimuli to be constructed.

Each stimulus was probed six times. The probes were constructed from the stimuli by adding "X" marks inside one or two of the gray blocks. The "X" marks were made of relatively thin lines (1 pixel wide) in order to force the subjects to attend to the stimuli carefully. If two "X" marks were present, they were in blocks that were 180° apart. There were an equal number of stimuli with one X probe and with two X probes, and the probes appeared with equal probability in each block location. Furthermore, no stimulus had only "yes" or only "no" probes, and the stimuli with large radii and the stimuli with small radii had equal numbers of "yes" and "no" probes. In this way, 60 unique trials were produced. Twelve stimuli were used in the practice trials, all of which had the squares at the 0°, 90°, 180°, and 270° positions. The remaining 48 stimuli were used in the test experiment and were ordered so that no two stimuli with the same configuration of blocks were in consecutive trials.

Procedure. As usual, a trial began with an exclamation point, which disappeared when the space bar was pressed. Following this, a set of brackets containing four gray blocks appeared and the subjects studied it. The subjects pressed the space bar and one or two "X" marks appeared inside the blocks. If two X marks appeared, the subjects were to respond "yes"; if only one X mark appeared, they were to respond "no".

Results and discussion

R.V. had a normal ability to attend to patterns subtending a relatively large visual angle, $t(7) = -1.32$, $p > .1$, for the response time scores, and $t(7) = 1.53$, $p > .1$, for the error rate scores. Thus, his earlier impaired performance cannot be ascribed, even in part, to a deficit in his ability to attend to larger regions of space.

Experiment XIV: Mental Rotation: Simultaneous Presentation

The results are painting a relatively complex picture, and it seemed worth collecting converging evidence from a different sort of task. The results from Experiments VI and VII showed that R.V. has difficulty encoding spatial relations, and the results from Experiment VIII showed that he has trouble storing information about location. If so, we reasoned, then he should also have trouble rotating visual mental images; this requires that one store information about the locations of parts of an object as one transforms their spatial relations. The rotation task we used is based on one devised by Shepard and Metzler (1971), and required the subjects to determine whether two shapes were identical or mirror reversals. The two figures were presented simultaneously, with left figure oriented vertically, and the right tilted one of 4 different angles; the top cell of both stimuli was filled in, making it easy to discover how the right figure was tilted.

The *manipulation* here was the amount of tilt. Shepard and Cooper (1982) review much data indicating that the greater the angular disparity between the stimuli in this task, the more "mental rotation" is required before they can be compared. These data can be explained if we posit a (very coarsely characterized) subsystem that shifts representations in the visual buffer. Kosslyn (1987) develops this idea in some detail, and argues that spatial relations encoding and property lookup subsystems must be involved in this process to use stored information to keep shapes properly aligned as they are being transformed. If so, then R.V. should have difficulty rotating objects in mental images. The *score* was the slope of the increase in times and errors with increasing angular disparity; the linear component of R.V.'s increase was compared to that of the control subjects.

Method

Materials. In this rotation task, both the unrotated standard shape and the rotated probe shape were presented simultaneously, side-by side. Cells of a 4 x 5 grid were chosen at random with the constraint that they form a single shape; the shapes had either two or three perceptual units (contiguous cells that form a bar). The remaining cells were eliminated, leaving only the shape. The top of each shape was marked by filling in a cell (black). Four shapes had two perceptual units, and four had three units. The single longest axes through the shape was oriented vertically or at 90, 135, or 180° clockwise rotations from the upright. The "no" trials included mirror-reversed shapes, whereas the "yes" trials included identical shapes. This experiment had two parts, both of which were counterbalanced by a Latin square design. Each part consisted of 32 trials. When taken together, the two halves were completely balanced for the different angles, responses, and stimulus complexities. There was also a balanced practice session, consisting of 16 trials using two target shapes that were not used in the actual experiment. This rotation task minimized the memory requirements needed to perform well.

Procedure. At the beginning of the trial an exclamation point appeared, and remained on the screen until the subjects pressed the space bar. The exclamation point disappeared and the screen went blank; 500 ms later a centered fixation point appeared and remained for another 500 ms. This was followed by the standard and rotated probe shapes, which appeared simultaneously. The standard always appeared in an upright position (the black box was orientated towards the top of the screen) and was always to the left of the fixation point; the probe appeared in one of four relative orientations and was always to the right of the fixation point. The subjects were asked to compare the probe to the target, and to respond "yes" if the two shapes were the same regardless of their relative orientations, and "no" if the probe was the mirror-image of the target. After the subjects responded, the exclamation point returned, signaling the start of the next trial.

Results and discussion

As is illustrated in Figure 16, R.V. required more time to rotate objects in images than the controls, $t(7) = 6.50$, $p < .01$, but there was no difference in the increase in errors with tilt, $t < .1$.

Thus, R.V. did in fact have difficulty rotating images, which is consistent with the finding that the frontal lobes are selectively activated during mental rotation (as measured by regional cerebral blood flow; Deutsch, Bourbon, Papanicolaou, & Eisenberg, 1988). However, these studies find

that it is the right frontal lobe that is selectively activated during tasks like this one, which is not consistent with the fact that R.V. has a left-frontal lesion. We worried that R.V.'s poor performance may not reflect rotation per se, but rather the effects of having two stimuli present at the same time, which may have overloaded his perceptual organization processes (in the preprocessing subsystem posited by Kosslyn et al., 1990). Thus, we conducted the following experiment.

Insert Figure 16 About Here

Experiment XV: Mental Rotation: Sequential Presentation

The simultaneous mental rotation task minimizes the importance of stored information, given that the subjects can look back and forth between the two stimuli to make the comparison (Just & Carpenter, 1979). However, if R.V. has difficulty encoding too much visual material at the same time, or has difficulty in part-for-part scanning, this should affect his ability to compare shapes in this task. Presenting stimuli sequentially reduces the amount of perceptual input and precludes the part-for-part comparison that is possible when two figures are present simultaneously. The results from the Ventral Shape Comparison task indicated that R.V. could perform sequential matches and could remember stimuli of comparable complexity. Thus, this experiment was like the Ventral Shape Comparison task, with the addition of a rotation component. It was like the previous one except that the stimuli were presented sequentially. Again, the *manipulation* was the angular disparity between the two figures, and the *score* was the increase in time and errors with increasing disparity.

Method

Materials. The standard and probe shapes were identical to those used in the simultaneous rotation experiment. The only change was that the target and probe shapes were now presented separately. The order of the trials was the same as for the simultaneous rotation experiment, and all other aspects of the design for the two experiments were identical.

Procedure. At the start of the trial, an exclamation point appeared for 500 ms. This was followed by a blank screen for 500 ms, after which the standard shape appeared in the center of the screen. The subjects were asked to study the shape for as long as they needed to memorize it. When ready, they pressed the space bar and the standard shape disappeared, and was replaced by a blank screen for 500 ms, after which the probe shape appeared. The subjects compared the probe shape to the standard shape stored in memory and responded "yes" if the shapes were the same, and "no" if they were mirror reversed.

Results and discussion

The results are illustrated in Figure 17. Again, R.V. rotated objects in images more slowly than did the control subjects, $t(7) = 18.20$, $p < .001$, but had essentially the same error score as the control subjects, $t(7) = 1.74$, $p > .1$. Thus, the deficit observed in the previous experiment cannot be ascribed solely to impaired scanning used in a part-for-part comparison process.

Insert Figure 17 About Here

Summary and Conclusions from the Case Study

The present case study demonstrates the utility of our theory of high-level vision as a guide to examining and interpreting dissociations in performance following brain damage. This investigation resulted in a profile of impaired and spared processing. Modern neuropsychology began with the study of dissociations between preserved and impaired abilities following brain damage (e.g., see Jackson, 1874). In many cases, however, the interpretation of such dissociations has been guided more by intuition than by a detailed theory of processing in the normal system. Indeed, only Marr's (1982) computational theory of visual processing has had an impact on the study and interpretation of visual deficits (e.g., Ratcliff, 1982; Riddoch & Humphreys, 1987). Marr's theory, however, did not provide a detailed decomposition of the structure of the higher-level visual processes, and cannot be used to guide precise examinations of patterns of deficits.

After we established a deficit in R.V.'s ability to compare two shapes presented sequentially in a grid, we used the theory to formulate a number of possible accounts for this deficit. Consider the status of each hypothesis in turn.

Viable hypotheses

The results allow us to rule out some of the hypotheses offered by the theory, and left others as plausible accounts of R.V.'s visual-spatial deficit.

1. *Visual buffer.* The visual buffer could have had regions of hypometabolism or scotoma, and the more complex the figure, the more likely it was to fall on a dysfunctional portion of the buffer. If this were the case, eliminating the grid lines should not have affected performance, but it did. Therefore, this hypothesis can be ruled out. Furthermore, there was no evidence of occipital dysfunction from PET.

2. *Attention window.* The attention window could have been restricted, so that only part of the figure could be seen at once. We eliminated this hypothesis directly.

3. *Preprocessing.* The preprocessing subsystem may have been impaired so that only a limited number of lines, symmetries, points of intersection, and other "nonaccidental properties" could be extracted at once. If so, then the presence of the grid lines may have overloaded the subsystem—forcing it to encode one part at a time. We have evidence that is consistent with this hypothesis.

4. *Pattern activation.* The pattern activation subsystem may have been damaged so that it could not store a representation of the first stimulus of the pair. We were able to rule out this possibility; when R.V. was asked to remember a pattern and later to decide whether an X would have fallen on it, there was no deficit. In addition, the pattern activation subsystem may have been damaged so that input from the preprocessing subsystem could not be compared properly to stored representations. We were able to eliminate this possibility by showing that R.V. could compare stimuli effectively when the grid lines were removed.

At first glance, we were puzzled about the apparent intact functioning of the pattern activation subsystem. Damage to the ventral system should have retarded the time to compare shapes, if nothing else. However, we must note that the damage was unilateral, and to the left side. Smith and Milner (1972) found that unilateral resection of the left temporal lobe did not affect memory for pictures, although resection of the right temporal lobe did. Given this finding, we then were led to ask why this damage affected the preprocessing subsystem? One possibility is suggested by PET scanning results summarized by Posner, Petersen, Fox, and Raichle (1988). Posner et al. found more activity in the left occipital temporal area when subjects saw words; it is possible that this area has been "tuned" to encode lines and angles during reading, and hence performance was impaired when grid lines were included. Anecdotally, it may be worth noting that although R.V. could read, he was very slow and awkward; prior to the stroke, he was an avid and fluent reader.

5. *Spatiotopic mapping.* The cells in the grid could have been encoded as separate locations, using the dorsal system. If so, then the impaired performance may reflect properties of the dorsal system. One possibility was that R.V.'s spatiotopic mapping subsystem was sensitive to the complexity of an object that shifts location, and establishing location requires more time for complex objects. This hypothesis was ruled out by showing that the deficit was present even when the stimulus was not displaced (in Experiment II), but it could be eliminated even when the stimulus was displaced (e.g., in Experiment III).

6. *Categorical spatial relations encoding.* If a pattern were encoded as a configuration of locations, they must have been specified relative to the grid itself. Categorical spatial relations (e.g., "top row, leftmost cell") are efficient for encoding such locations. We showed that R.V. had a deficit in encoding at least one categorical spatial relation, above/below. Thus, if he were using this subsystem to encode locations of filled cells, more time may have been required to encode the more complex patterns.

7. *Coordinate spatial relations encoding.* The locations of the filled cells could be specified using metric distances. We also found that R.V. did in fact have a deficit in encoding metric spatial relations.

8. *Associative memory.* The input from either the dorsal or ventral system (or both) may not have been reliably sent to associative memory, prior to reaching a judgment. We found that R.V. had difficulty storing location information, which may have reflected the damage to dorsolateral prefrontal cortex. However, we have no evidence that R.V. had trouble using visual-spatial information once it was encoded into long-term memory.

9. *Top-down processing.* Because the "different" stimuli were relatively similar to the ones studied initially, "second looks" may have been used at least some of the time. It is possible that such processing is used more often with more complex stimuli because they are more difficult to represent fully in a single encoding. We found that R.V. had no difficulty accessing information about location or shape stored in long-term memory, and using such information to direct attention.

10. *Attention shifting.* Even the control subjects may have examined the stimuli a part at a time, but it was possible that they were able to shift their attention (i.e., scan over it) much faster than R.V. We had inconclusive results here, with a speed-accuracy tradeoff making it difficult to draw firm conclusions.

Finally, we also found that R.V. had a deficit in mentally rotating objects. This deficit is consistent with our finding that he had trouble representing spatial location, given that one must store such information as one mentally manipulates the orientations of the patterns.

The present approach, then, is a departure from the usual technique in neuropsychology of establishing pairs of dissociations and associations following brain damage (e.g., Caramazza, 1986). We recognize that lesions are often relatively large, and sometimes have remote effects by denervating other parts of the brain. In this research we found evidence of a system of functional impairments, which appears to reflect dysfunction in the occipital-temporal junction area (which putatively implements the preprocessing subsystem) and the frontal lobes (which are critically involved in encoding and storing spatial information). This approach is admittedly more complex than the usual fare in neuropsychology, but seems fitting for a description of the dysfunction of a marvelously complex organ, the brain.

Footnotes

Footnote 1. I owe the idea that the magnocellular ganglia may project preferentially to the right hemisphere to Marge Livingstone.

Footnote 2. One can also ask why the right hemisphere monitors larger fields than the left rather than vice versa. A possible account rests on three ideas. First, the right hemisphere is more mature at birth (Taylor, 1969). Second, the infant, having little information in memory to guide attention, relies heavily on preattentive processes in vision. These processes are more effective if large receptive fields are monitored. Third, once the right hemisphere has been used heavily for this purpose, considerable neural reconfiguration would be required to allow it to be effective in controlling focal attention mechanisms. Hence, when the left hemisphere matures, it is able to accomplish these tasks easier than the right, and the specialization develops. (This idea was inspired by those of de Schonen and Mathivet, 1989; Hellige, 1989; and Sargent, 1988).

Footnote 3. Kosslyn et al. (1990) pointed out that because categorical spatial relations do not specify precise positions, additional processes are necessary to convert such representations to specific locations in a given image. They posited a separate subsystem to perform these conversions. I am no longer certain that this distinction is justified, and will be conservative by assuming for the moment that the categorical property lookup subsystem may perform the necessary conversion by itself.

Footnote 4. In either case, one cannot use the position information to adjust directly the location of an image in the visual buffer, without first moving the attention window; spatial relations are always specified relative to some part of an object or scene, and so the size and orientation of the object or scene will determine where the new part belongs. And the size and orientation of the object or scene is only explicit in the visual buffer, and may vary from instance to instance.

References

- Allman, J. M., and Kaas, J. H. (1976). Representation of the visual field on the medial wall of occipital-parietal cortex in the owl monkey. *Science*, 191, 572-575.
- Andersen, R. A. (1987). The role of the inferior parietal lobule in spatial perception and visual-motor integration. In F. Plum, V. B. Mountcastle, and S. R. Geiger (Eds.), *Handbook of Physiology: The Nervous System V, Higher Functions of the Brain Part 2*. Bethesda, MD: American Physiological Society.
- Andersen, R. A., Essick, G. K., and Siegel, R. M. (1985). Encoding of spatial location by posterior parietal neurons. *Science*, 230, 456-458.
- Attneave, F. (1974). How do you know? *American Psychologist*, 29, 493-499.
- Baddeley, A. (1976). *Working Memory*. Oxford: Oxford University Press.
- Bailey, P., and von Bonin, G. (1951). *The Isocortex of Man*. Urbana, IL: University of Illinois Press.
- Bauer, R. M., and Rubens, A. B. (1985). Agnosia. In K. M. Heilman and E. Valenstein (Eds.), *Clinical Neuropsychology*. New York: Oxford University Press.
- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94, 115-147.
- Biederman, I., and Shiffrar, M. M. (1987). Sexing day-old chicks: a case study and expert systems analysis of a difficult perceptual-learning task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13, 640-645.
- Bower, G. H., and Glass, A. L. (1976). Structural units and the reintegrative power of picture fragments. *Journal of Experimental Psychology: Human Learning and Memory*, 2, 456-466.
- Broadmann, K. (1909). *Vergleichende Lokalisationslehre der Grosshirnrinde*. Leipzig: Barth.
- Campion, J. (1987). Apperceptive agnosia: the specification and description of constructs. In Humphreys, G. W., and Riddoch, M. J. (1987a) (Eds.), *Visual Object Processing: A Cognitive Neuropsychological Approach*. Hillsdale, NJ: Erlbaum. pp. 197 - 232.
- Cave, K. R., and Kosslyn, S. M. (1989). Varieties of size-specific visual selection. *Journal of Experimental Psychology: General*, 118, 148-164.
- Cohen, G. (1975). Hemispheric differences in the utilization of advance information. In P. M. A. Rabbit and S. Dornic (Eds.), *Attention and Performance V*. New York: Academic Press.
- Cooper, L. A., and Shepard, R. N. (1973). Chronometric studies of the rotation of mental images. In W. G. Chase (Ed.), *Visual Information Processing*. New York: Academic Press.
- Corbetta, M., Miezin, F. M., Dobmeyer, S., Shulman, G. L., and Petersen, S. E. (1990). Attentional modulation of neural processing of shape, color, and velocity in humans. *Science*, 248, 1556-1559.
- Cowey, A. (1985). Aspects of cortical organization related to selective attention and selective impairments of visual perception: a tutorial review. In M. S. Posner and O. S. Marin (Eds.), *Attention and Performance XI*. Hillsdale, NJ: Erlbaum.
- Critchley, M. (1953). *The Parietal Lobes*. Edward Arnold: London.
- Cutting, J. E., and Kozlowski, L. T. (1977). Recognizing friends by their walk: Gait perception without familiarity. *Bulletin of the Psychonomic Society*, 9, 353-356.
- Cutting, J. E., and Proffitt, D. R. (1981). Gait perception as an example of how we may perceive events. In R. D. Walk & H. L. Pick (Eds.), *Intersensory Perception and Sensory Integration*. New York: Plenum Press.
- Damasio, A. R. (1986). Disorders of complex visual processing: agnosias, achromatopsia, Balint's syndrome and related difficulties of orientation and construction. In M.-M. Mesulam (Ed.), *Principles of Behavioral Neurology*. Philadelphia, PA: F. A. Davis. (pp 259-288)
- Damasio, A. R., Damasio, H., and Van Hoesen, G. W. (1982). Prosopagnosia: Anatomic basis and behavioral mechanisms. *Neurology (NY)*, 32, 331-341.
- De Renzi, E. (1982). *Disorders of Space Exploration and Cognition*. New York: John Wiley & Sons.
- de Schonen, S., and Mathivet, E. (1989). First come, first served: A scenario about the development of hemispheric specialization in face recognition during infancy. *European Bulletin of Cognitive Psychology*, 1, 3-44.

- Delis, D. C., Robertson, L. C., and Efron, R. (1986). Hemispheric specialization of memory for visual hierarchical stimuli. *Neuropsychologia*, 24, 205-214.
- Deutsch, G., Bourbon, W. T., Papanicolaou, A. C., & Eisenberg, H. M. (1988). Visuospatial experiments compared via activation of regional cerebral blood flow. *Neuropsychologia*, 26, 445-452.
- Farah, M. J. (1988). Is visual imagery really visual? Overlooked evidence from neuropsychology. *Psychological Review*, 95, 307 - 317.
- Farah, M. J., and Hammond, K. M. (1988). Mental rotation and orientation-invariant object recognition: dissociable processes. *Cognition*, 29, 29 - 46.
- Felleman, D. J., and Van Essen, D. C. (in press). Distributed hierarchical processing in primate cerebral cortex. *Cerebral Cortex*.
- Fisk, J. D., and Goodale, M. A. (1988). The effects of unilateral brain damage on visually guided reaching: hemispheric differences in the nature of the deficit. *Experimental Brain Research*, 72, 425-435.
- Frackowiak, R.S., Lenzi, G.L., Jones, T., and Heather, J.D. (1980). Quantitative measurement of regional cerebral blood flow and oxygen metabolism in man using ¹⁵O and positron emission tomography: Theory, procedure and normal values. *Journal of Computer Assisted Tomography*, 4, 727-726.
- Goldman-Rakic, P. S. (1987). Circuitry of primate prefrontal cortex and regulation of behavior by representational knowledge. In F. Plum and V. Mountcastle (Eds.), *Handbook of Physiology*, Vol 5. Washington, D.C.: American Physiological Society. (pp. 373-417)
- Goodale, M. A. (1988). Hemispheric differences in motor control. *Behavioural Brain Research*, 30, 203-214.
- Gregory, R. L. (1970). *The Intelligent Eye*. London: Weidenfeld and Nicholson.
- Gross, C. G., Desimone, R., Albright, T. D., and Schwartz, E. L. (1984). Inferior temporal cortex as a visual integration area. In F. Reinoso-Suarez & Ajmonc-Marsan (Eds.), *Cortical Integration*. New York: Raven Press, pp 291-315.
- Hannay, H. J., Varney, N. R., and Benton, A. L. (1976). Visual localization in patients with unilateral brain disease. *Journal of Neurology, Neurosurgery and Psychiatry*, 39, 307-313.
- Harries, M. H., and Perrett, D. I. (in press). Visual processing of faces in temporal cortex: physiological evidence for a modular organization and possible anatomical correlates. *Journal of Cognitive Neuroscience*.
- Hebb, D. O. (1949). *The Organization of Behavior*. New York: Wiley.
- Hecaen, H., and Albert, M. L. (1978). *Human Neuropsychology*. New York: John Wiley and Sons.
- Hellige, J. B. (1989). Endogenous and experimental determinants of cerebral laterality: What develops? *European Bulletin of Cognitive Psychology*, 1, 85-89.
- Hellige, J. B. and Michimata, C. (1989). Categorization versus distance: Hemispheric differences for processing spatial information. *Memory & Cognition*, 17, 770-776.
- Hock, H., Kronseder, C., and Sissons, S. (1981). Hemispheric asymmetry: the effect of orientation on same different comparison. *Neuropsychologia*, 19, 723-727.
- Holmes, G. (1919). Disturbances of visual space perception. *British Medical Journal*, 2, 230-233.
- Holmes, E.J., and Gross, C. G. (1984). Effects of temporal lesions on discrimination of stimuli differing in orientation. *Journal of Neuroscience*, 4, 3063-3068.
- Humphreys, G. W., and Riddoch, M. J. (1987a) (Eds.). *Visual Object Processing: A Cognitive Neuropsychological Approach*. Hillsdale, NJ: Erlbaum.
- Humphreys, G. W., and Riddoch, M. J. (1987b). *To See But Not To See: A Case Study of Visual Agnosia*. Hillsdale, NJ: Erlbaum.
- Hyvarinen, J. (1982). Posterior parietal lobe of the primate brain. *Physiological Reviews*, 62, 1060-1129.
- Jackson, J. H. (1874). On the duality of the brain. Reprinted (1932) in J. Taylor (Ed.), *Selected Writings of John Hughlings Jackson*. London: Hodder and Stoughton.
- Johansson, G. (1950). *Configurations in Event Perception*. Uppsala, Sweden: Almqvist & Wiksell.
- Johansson, G. (1975). Visual motion perception. *Scientific American*, 232, 76-88.

- Jolicoeur, P. (in press). On the role of mental rotation and feature extraction in the identification of disoriented objects: a dual-systems theory. *Mind and Language*,
- Kinsbourne, M., and Warrington, E. K. (1963). The localizing significance of limited simultaneous visual form perception. *Brain*, 86, 697-702.
- Kirby, K.N., and Kosslyn, S. M. (in press). Thinking visually. *Mind and Language*,
- Koenig, O., Kosslyn, S. M., and Chabris, C. F., and Gabrieli, J. D. E.. (1990). Computational constraints on the acquisition of spatial knowledge in the cerebral hemispheres. University of Geneva manuscript.
- Koenig, O., Reiss, L., and Kosslyn, S. M. (in press). The development of categorical and coordinate spatial relations representations. *Child Development*,
- Kolb, B. and Whishaw, I. Q. (1985). *Fundamentals of Human Neuropsychology* (2nd Ed.). San Francisco: W. H. Freeman.
- Kosslyn, S. M. (1980). *Image and Mind*. Cambridge, MA: Harvard University Press.
- Kosslyn, S. M. (1987). Seeing and imagining in the cerebral hemispheres: a computational approach. *Psychological Review*, 94, 148-175.
- Kosslyn, S. M. (1988). Aspects of a cognitive neuroscience of mental imagery. *Science*, 240, 1621-1626.
- Kosslyn, S. M. (in press). A cognitive neuroscience of visual mental imagery: further developments. In R. Logie (Ed.), *Advances in Mental Imagery Research*. Hillsdale, NJ: Erlbaum.
- Kosslyn, S. M., Berndt, R. S., and Doyle, T. J. (1985). Imagery and language: a preliminary neuropsychological investigation. In M. S. Posner and O. S. Marin (Eds.), *Attention and Performance XI*. Hillsdale, NJ: Erlbaum.
- Kosslyn, S. M., Cave, C. B., Provost, D., and Von Gierke, S. (1988). Sequential processes in image generation. *Cognitive Psychology*, 20, 319-343.
- Kosslyn, S. M., Chabris, C. F., , and Marsolek, C. J., and Koenig, O. (in press). Categorical versus coordinate spatial representations: Computational analyses and computer simulations. *Journal of Experimental Psychology: Human Perception and Performance*,
- Kosslyn, S. M., Daly, P., McPeck, R., Alpert, N. A., and Caviness, V. S. (1990). Systems of dysfunction: A theory-driven analysis of a deficit in high-level visual processes. Harvard University manuscript.
- Kosslyn, S. M., Flynn, R. A., Amsterdam, J. B., and Wang, G. (1990). Components of high-level vision: A cognitive neuroscience analysis and accounts of neurological syndromes. *Cognition*, 34, 203-277.
- Kosslyn, S. M., Hillger, L. A., Engel, S. E., Clegg, B., and Hamilton, S. E. (1990). Attention-based and form-based imagery: Evidence from studies of transfer of practice. Harvard University manuscript.
- Kosslyn, S. M., Hillger, L. A., Livingstone, M. S., and Hamilton, S. E. (1990). Does the right hemisphere monitor larger receptive fields than the left? Harvard University manuscript.
- Kosslyn, S. M., Koenig, O., Barrett, A., Cave, C. B., Tang, J., and Gabrieli, J. D. E. (1989). Evidence for two types of spatial representations: Hemispheric specialization for categorical and coordinate relations. *Journal of Experimental Psychology: Human Perception and Performance*, 15, 723-735.
- Kosslyn, S. M., and Park, S. (1990). Incidental memory for left-right orientation in the cerebral hemispheres. Harvard University manuscript.
- Kosslyn, S. M., Segar, C., Pani, J., and Hillger, L. A. (in press). When is imagery used? A diary study. *Journal of Mental Imagery*.
- Krieg, W. J. S. (1973). *Architectonics of the Human Cerebral Fiber Systems*. Evanston, ILL: Brain Books.
- Larsen, A., and Bundesen, C. (1978). Size scaling in visual pattern recognition. *Journal of Experimental Psychology: Human Perception and Performance*, 4, 1 - 20.
- Levine, D. N. (1982). Visual agnosia in monkey and man. In D. J. Ingle, M.A. Goodale, and R. J. W. Mansfield (Eds.), *Analysis of Visual Behavior*. Cambridge, MA: MIT Press.
- Levine, D. N., Maini, R. B., and Calvanio, R. (1988). Pure agraphia and Gerstmann's Syndrome as a visuospatial-language dissociation: An experimental case study. *Brain and Language*, 35, 172-196.
- Litton, J., and Bergstrom, M., and Eriksson, L., et al. (1984). Performance study of the PC-384 PET camera for the brain. *Journal of Computer Assisted Tomography*, 8, 74-87

- Livingstone, M. S., and Hubel, D. H. (1987). Psychophysical evidence for separate channels for the perception of form, color, movement and depth. *Journal of Neuroscience*, 7, 3416-3468.
- Lowe, D. G. (1987a). Three-dimensional object recognition from single two-dimensional images. *Artificial Intelligence*, 31, 355-395.
- Lowe, D. G. (1987b). The viewpoint consistency constraint. *International Journal of Computer Vision*, 1, 57-72.
- Luria, A. R. (1959). Disorders of "simultaneous perception" in a case of bilateral occipito-parietal brain injury. *Brain*, 82, 437-449.
- Luria, A. R. (1980). *Higher Cortical Functions in Man*. New York: Basic Books.
- Marr, D. (1982). *Vision*. San Francisco, CA: W. H. Freeman.
- Maunsell, J. H. R., and Newsome, W. T. (1987). Visual processing in monkey extrastriate cortex. *Annual Review of Neuroscience*, 10, 363-401.
- McLeod, P., McLaughlin, C., and Nimmo-Smith, I. (1985). Information encapsulation and automaticity: evidence from the visual control of finely timed actions. In M.I. Posner & O.S.M. Marin (Eds.), *Attention and Performance XI*. Hillsdale, NJ: Erlbaum. pp 391-406.
- Mehta, Z., Newcombe, F., and Damasio, H. (1987). A left hemisphere contribution to visuospatial processing. *Cortex*, 23, 447-461.
- Mesulam, M-M. (1981). A cortical network for directed attention and unilateral neglect. *Annals of Neurology*, 10, 309-325.
- Mishkin, M., Ungerleider, L. G., and Macko, K. A. (1983). Object vision and spatial vision: Two cortical pathways. *Trends in NeuroSciences*, 6, 414-417.
- Moran, J., and Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science*, 229, 782-784.
- Navon, D., and Norman, J. (1983). Does global precedence really depend on visual angle? *Journal of Experimental Psychology: Human Perception and Performance*, 9, 955-965.
- Nickerson R.S., & Adams, M.J. (1976). Long term memory for a common object. *Cognitive Psychology*, 11, 287-307.
- Olson, D., and Bialystok, E. (1983). *Spatial Cognition: the Structure and Development of Mental Representations of Spatial Relations*. Hillsdale, NJ: Erlbaum.
- Paivio, A. (1971). *Imagery and Verbal Processes*. New York: Holt, Rinehart & Winston.
- Palmer, S. E. (1977). Hierarchical structure in perceptual representations. *Cognitive Psychology*, 9, 441-474.
- Pandya, D. N., and Yeterian, E. H. (1985). Architecture and connection of cortical association areas. In A. Peters & E. G. Jones (Eds.), *Cerebral Cortex, Vol 4: Association and Auditory Cortices*. New York: Plenum.
- Perrett, D. I., Smith, P. A. J., Potter, D. D., Mistlin, A. J., Head, A. S., Milner, A. D., and Jeeves, M. A. (1985). Visual cells in the temporal cortex sensitive to face view and gaze direction. *Proceedings of the Royal Society London, B* 223, 293-317.
- Podgorny, P., and Shepard, R. N. (1978). Functional representations common to visual perception and imagination. *Journal of Experimental Psychology: Human Perception and Performance*, 4, 21-35.
- Pohl, W. (1973). Dissociation of spatial discrimination deficits following frontal and parietal lesions in monkeys. *Journal of Comparative and Physiological Psychology*, 82, 227-239.
- Posner, M. I., Inhoff, A. W., Friedrich, F. J., and Cohen, A. (1987). Isolating attentional systems: A cognitive-anatomical analysis. *Psychobiology*, 15, 107 - 121.
- Posner, M. I., Petersen, S. E., Fox, P. T., and Raichle, M. E. (1988). Localization of cognitive operations in the human brain. *Science*, 240, 1627-1631.
- Posner, M. I., Walker, J. A., Friedrich, F. J., and Rafal, R. D. (1984). Effects of parietal lobe injury on covert orienting of visual attention. *Journal of Neuroscience*, 4, 1863-1974.
- Ratcliff, G. (1982). Disturbances of spatial orientation associated with cerebral lesions. In M. Potegal (Ed.), *Spatial Abilities: Developmental and Physiological Foundations*. New York: Academic Press.
- Ratcliff, G., and Davies-Jones, G. A. G. (1972). Defective visual localisation in focal brain wounds. *Brain*, 95, 49-60.

- Reed, S. K., and Johnsen, J. A. (1975). Detection of parts in patterns and images. *Memory and Cognition*, 3, 569-575.
- Rock, I. (1973). *Orientation and Form*. New York: Academic Press.
- Senda, M., Buxton, R. B., Alpert, N. M., and Correia, J. A., Mackay, B. C., Weise, S. B., and Ackerman, R. H. (1988). The ¹⁵O steady-state method: Correction for variation in arterial concentration, *Journal of Cerebral Blood Flow and Metabolism*, 8, 681-690.
- Sergent, J. (1988). Face perception and the right hemisphere. In L. Weiskrantz (Ed.), *Thought without Language*. Oxford, UK: Oxford University Press. pp 108-131.
- Sergent, J. (in press). title. *Journal of Experimental Psychology: Human Perception and Performance*.
- Shepard, R. N., & Cooper, L. A. (1982). *Mental Images and their Transformations*. Cambridge, MA: MIT Press.
- Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs*, 74, (11, Whole no. 498).
- Squire, L. R. (1987). *Memory and Brain*. New York: Oxford University Press.
- Talairach, J., and Sziklag, L. (1967). *Atlas of Stereotactic Anatomy of the Telencephalon*. Paris: Masson and Cie.
- Tarr, M. J., & Pinker, S. (1989). Mental rotation and orientation dependence in shape recognition. *Cognitive Psychology*, 21, 233-282.
- Taylor, D. C. (1969). Differential rates of cerebral maturation between sexes and between hemispheres. *Lancet*, 2, 140-148.
- Taylor, A. M., and Warrington, E. K. (1973). Visual discrimination in patients with localized brain lesions. *Cortex*, 9, 82-93.
- Treisman, A. M., and Gelade, G. (1980). A feature integration theory of attention. *Cognitive Psychology*, 12, 97-136.
- Tyler, H. R. (1968). Abnormalities of perception with defective eye movements (Balint's syndrome). *Cortex*, 4, 154-171.
- Ullman, S. (1979). *The Interpretation of Visual Motion*. Cambridge, MA: MIT Press.
- Ullman, S. (1989). Aligning pictorial descriptions: An approach to object recognition. *Cognition*, 32, 193 - 254.
- Ungerleider, L. G., and Mishkin, M. (1982). Two cortical visual systems. In D. J. Ingle, M. A. Goodale, and R. J. W. Mansfield (Eds.), *Analysis of Visual Behavior*. Cambridge, MA: MIT Press.
- Van Essen, D. (1985). Functional organization of primate visual cortex. In A. Peters and E. G. Jones (Eds.), *Cerebral Cortex (Vol. 3)*. New York: Plenum Press.
- Van Essen, D. C., and Maunsell, J. H. R. (1983). Hierarchical organization and functional streams in visual cortex. *Trends in Neuroscience*, September, 370 - 375.
- Van Essen, D. C., Felleman, D. J., DeYoe, E. A., Olavarria, J., and Knierim, J. (1990). Modular and hierarchical organization of extrastriate visual cortex in the macaque monkey. *Quantitative Biology*, 55,
- Van Kleeck, M. H. (1989). Hemispheric differences in global versus local processing of hierarchical visual stimuli by normal subjects: New data and a meta-analysis of previous studies. *Neuropsychologia*, 27, 1165-1178.
- Warrington, E. K., and James, M. (1967). Disorders of visual perception in patients with localized cerebral lesions. *Neuropsychologia*, 5, 253-266.
- Warrington, E. K., and Rabin, P. (1970). Perceptual matching in patients with cerebral lesions. *Neuropsychologia*, 8, 475-487.
- Williams, M. (1970). *Brain Damage and the Mind*. Penguin Book: Middlesex, England.
- Yarbus (1967). *Eye Movements and Vision*. New York: Plenum Press.
- Yin, T.C.T., and Mountcastle, V. B. (1977). Visual input to the visuomotor mechanisms of the monkey's parietal lobe. *Science*, 197, 1381-1383.
- Zipser, D. and Andersen, R. A. (1988). A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature*, 331, 679-684.

Figures

- Figure 1. The subsystems of high-level vision inferred by Kosslyn et al. (1990) and modified by Kosslyn (in press). The heavy black lines group subsystems into sets, as described in the text.
- Figure 2. Areas of damage, as determined by MRI and PET scan.
- Figure 3. The trial sequence for the Shape Comparison experiment.
- Figure 4. Results from the Shape Comparison experiment.
- Figure 5. Results from the Short-Term Memory Control experiment.
- Figure 6. The trial sequence for the Pattern Activation Encoding experiment.
- Figure 7. Results from the Pattern Activation Encoding experiment.
- Figure 8. Results from the Preprocessing Overload experiment.
- Figure 9. Results from the Ventral Shape Comparison experiment.
- Figure 10. The trial sequence from the Categorical Spatial Relations Encoding experiment.
- Figure 11. Results from the Location Associative Memory experiment.
- Figure 12. Results from the Preprocessing Followup experiment.
- Figure 13. The trial sequence for the Shape Top-Down Search experiment.
- Figure 14. The trial sequence for the Scanning experiment.
- Figure 15. Results from the Scanning experiment.
- Figure 16. Results from the Mental Rotation: Simultaneous Presentation experiment.
- Figure 17. Results from the Mental Rotation: Sequential Presentation experiment.





fig 1

MRI RESULTS



PET RESULTS







Fig 2




(warning)



(study)



(3s delay)



(test stimulus)




Fig.4 Results from
Shape Comparison




Fig. 5 - Results from
Shape Short-Term Memory
Control




Fig 6




Fig. 7 - Results from
Pattern Activation Encoding




Fig 8 - Results From
Preprocessor Overload







Fig 9- Ventral Shape
Comparison



(warning)



(1s delay)



(test stimulus)

Fig 10




Fig. 11 - Location
Associative Memory




Fig. 12 -
Preprocessor Followup




fig 13




fig 14




Fig. 15 - Results from Scanning




Fig. 16 - Mental Rotation:
Simultaneous Presentation




Fig. 17- Mental Rotation;
Sequential Presentation